2009 COMPUTATION

Annual Report

Lawrence Livermore National Laboratory

ANTICIPATE · INNOVATE · DELIVER

About the Cover

Cover design by Daniel Moore

A 400-million-particle molecular dynamics simulation, run on the Sequoia Initial Delivery System (Dawn), was used as the basis for the artistic abstraction featured on the cover.

Computation Directorate

NNSA National Nuclear Security Administration

LAWRENCE LIVERMORE NATIONAL SECURITY LLC

Production team

LLNL Associate Director for Computation

Dona Crawford

Deputy Associate Directors

Ronald Schmucker (Principal Deputy), David L. Brown, and Michel McCoy

Scientific editors

Benjamin Grover and Tyler J. Alumbaugh

Art Director

Daniel Moore

Production Editor

Deanna Willis

Proofreader

Pamela MacGregor

Photographers

Nancy Rutter and Jacquelin McBride

Print production

Diana Horne and Monarch Print Copy and Design Solutions

LLNL-TR-427019

Contents

■ Anticipate ■ Innovate ■ Deliver

A Message from the Associate Director

Dona Crawford
Associate Director
Computation Directorate

In a January 2010 editorial in *The Wall Street Journal*, Vice President Joe Biden cast a spotlight on the outstanding scientists and engineers employed at our national laboratories. "The same [National Lab] experts who maintain our [nuclear] arsenal play a key role in guaranteeing our country's security now and for the future," Biden said. He also recognized the leading role of the laboratories in "develop[ing] tomorrow's cutting-edge technologies for our security and prosperity." High-performance computing is essential to this technology development, and the Computation Directorate at Lawrence Livermore National Laboratory (LLNL) stands at the heart of that enterprise. Our mission is to anticipate the needs of the country, develop innovative solutions to pressing national and global problems, and deliver solutions to new challenges in order to keep our country, economy, and people safe, secure, and prosperous.

This year's annual report highlights many examples of anticipation, innovation, and delivery by Computation's exceptional employees. Whether we are putting new collaboration and productivity tools into the hands of the user, writing the control software for the National Ignition Facility (NIF), developing mathematical models and algorithms for solving key scientific questions, or creating new platforms for high-performance computing, Computation is foundational to the Laboratory's success. Our solutions address computation and communication needs, from hand-held mobile devices for an increasingly connected workforce to machines capable of 500 trillion floating operations per second (teraFLOPS).

A new 500-teraFLOPS Advanced Simulation and Computing system, called Dawn, is generating exciting results across a broad spectrum of applications, including climate modeling, laser–plasma interaction, and hydrodynamics. Dawn's initial results demonstrated groundbreaking science, enhanced code performance, and some of the largest, highest resolution simulations ever run in their respective scientific field. Dawn (an IBM BlueGene/P system) will lay the applications foundation for Sequoia, a 20-petaFLOPS IBM system due in 2012.

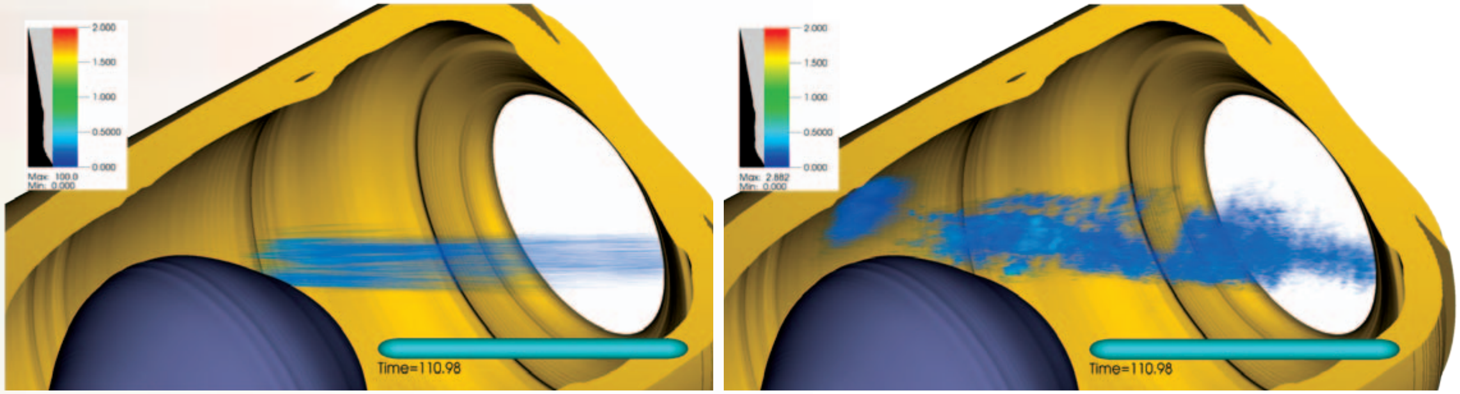

A laser beam propagation simulation, run on Dawn, shows simulated backscatter inside an ignition target (left: laser light resonantly scatters off ion acoustic waves; right: laser light resonantly scatters off electron plasma waves). Such simulations help optimize NIF target performance.

While nuclear weapons security remains a principal mission, the Laboratory is applying its expertise and capabilities to a broad range of national challenges. Through partnerships with the programs, Computation has delivered capabilities for space situational awareness and target system diagnostics for NIF. We have also delivered security tools for the Laboratory computer systems, software applications that improve nuclear weapon disassembly and inspection efficiencies, and a "Blue" network that facilitates international collaboration while ensuring security by limiting foreign nationals' access to only the assets they need.

Computational innovation has allowed scientific advances such as high-performance simulations of charged particles, biological theories for the evolution of viruses, and theories and models for laser–plasma interactions. Also, simulation codes have been improved by adding adaptive mesh refinement and improved linear and nonlinear solvers. Innovative thinking on the systems side has led to improved hardware and software for storage-intensive supercomputing, improved high-performance computing techniques for data streaming, and new approaches to cyber security.

By staying in tune with the needs of Laboratory programs, we also anticipate and confront tomorrow's challenges. Our track record of innovation will help us move into the exascale regime, a computational space wherein we are just beginning to unravel the challenges: constrained total energy, unreliability of architectures with very large numbers of components, and limitations in memory size and access speed, to name a few. Computation scientists will find ways to adapt and scale algorithms, improve and build new multiphysics simulation codes, and quantify the uncertainty in massive exascale calculations.

Our organization has been navigating uncharted computational waters for the nearly 60 years since Livermore's founding as a second nuclear weapons development laboratory. LLNL designed and developed many of the nuclear weapons in the nation's arsenal and now plays a critical role in the Stockpile Stewardship Program's mission to steward those weapons. A cornerstone of that effort is the Advanced Simulation and Computing (ASC) Program, which provides the computational tools needed to ensure the safety and reliability of the weapons without nuclear testing. This effort has required an unprecedented acceleration of computing capabilities made possible by collaboration with industry. In a trail-blazing partnership with IBM, we developed the BlueGene/L-class machines, systems that were of equal benefit to the Laboratory and IBM. The more recent Hyperion Project takes that successful partnership model one step further, bringing together multiple industry partners to establish an environment to test new hardware and software technologies for future systems, such as the 20-quadrillion-floating-operations-per-second (petaFLOPS) Sequoia platform due in 2012. Together with Sandia National Laboratories, LLNL is considering the creation of a Livermore Valley Open Campus (LVOC), a vehicle to promote and enhance collaboration between the laboratories, industry, and academia. Although the open campus concept is in its infancy, Computation is already planning for the open and agile collaborations an LVOC would make possible. These symbiotic experiences underscore LLNL's successful role as a trusted partner in developing unique solutions to difficult problems.

As Laboratory employees, we operate in a challenging environment that requires us to protect our machines, networks, and data from a constant barrage of cyber attacks, while expanding vital collaborations with industry and academia. To address these competing constraints, Computation initiated research in automated cyber security defenses and enacted methods and policies that permit openness in the Laboratory infrastructure while maintaining rigorous security measures. We also implemented more efficient procedures for approving new software and hardware. These techniques allow a quicker deployment to our scientists and engineers while maintaining the security posture necessary for a national laboratory. We upgraded the entire institution to Microsoft Office 2007/2008 over a single weekend using an automated technique that sets a new standard for other such software upgrades at the Laboratory. We continually explore options that keep our networks,

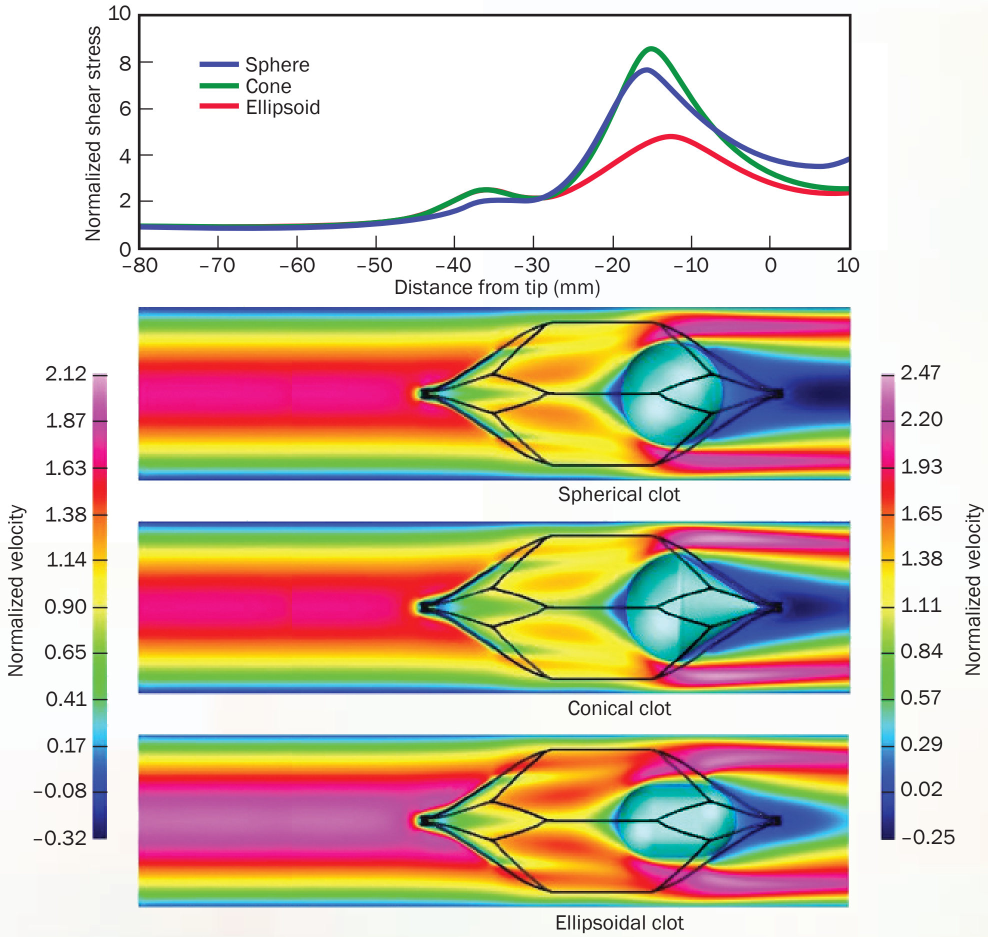

Computation researchers have developed simulation capabilities to study blood flow through a TrapEase inferior vena cava filter. The three-dimensional computational fluid dynamics simulations use Computation's Overture software framework to model blood flow past an unobstructed filter and a filter partially obstructed by blood clots of varying shapes and sizes. The computer modeling approach offers numerous advantages over traditional in-vitro techniques, including higher spatial resolution, easier adaptation for new filter designs and clot morphologies, and straightforward modification of flow parameters.

machines, and data secure while allowing our increasingly mobile workforce, collaborators, and vendors to take full advantage of rapidly evolving business communication technology.

The capabilities developed by Computation are integral to the advancement of the science basic to Department of Energy (DOE) missions. Computation provides critical support to NIF, the laser fusion experiment that began operations in 2009. Laboratory scientists were finalists for the 2009 Gordon Bell Prize for using parallel computing to simulate long-range electrostatic interactions in hot plasmas. The team's work not only "kicked the tires" on newer platforms but also allowed for higher fidelity NIF simulations. As the world's premier experimental laser facility, NIF will help pave the way for cleaner energy and better understanding of physical phenomena at temperatures and pressures like those found on the sun. Clearly, with energy security increasing in national priority, we need alternative sources of energy to keep the U.S. competitive and secure. Employees and programs throughout the directorate and the Laboratory are lending their expertise and ideas to help meet these challenges.

Computation not only provides scientific solutions for our primary sponsor, the National Nuclear Security Administration (NNSA), but also informs NNSA leaders' business decisions at the micro and macro levels. Our capabilities and resources have been used to help NNSA model its entire enterprise in order to provide a big-picture view of how and where to dedicate resources and a better understanding of funding trade-offs.

As we look to the future, many complex issues loom on the horizon, including national security challenges, energy security, environmental security, economic security, and climate change. I am confident that Computation will leverage its unique capabilities to anticipate, innovate, and deliver solutions that address these and other pressing national and global needs. I agree wholeheartedly with Vice President Biden when he said that the national laboratories and the experts they employ "are a national treasure that we must and will sustain." Computation continues to answer that high calling.

DONA CRAWFORD
ASSOCIATE DIRECTOR, COMPUTATION

RETURN TO TABLE OF CONTENTS

An Award-Winning Organization

Many Computation personnel and projects received prestigious external recognition in 2009. Some of the accomplishments are included in this section. The people who perform this award-winning work share several characteristics—dedication, tenacity, and passion—and represent the very best Computation has to offer.

Top Technological Innovation

A team of computer scientists won an R&D 100 Award for developing ROSE, software that makes compiler technology more accessible to programmers. The R&D 100 Awards, often dubbed the "Oscars of invention," recognize the most innovative ideas of the year and are a mark of excellence known to industry, government, and academia.

ROSE radically changes the accessibility of compiler technologies, allowing access to average software developers and scientists. ROSE enables users to build their own tools, including defect detection tools to uncover undetected bugs, code optimization tools to maximize program performance, and program transformation tools that allow users to easily develop programs for today's fast-changing hardware platforms.

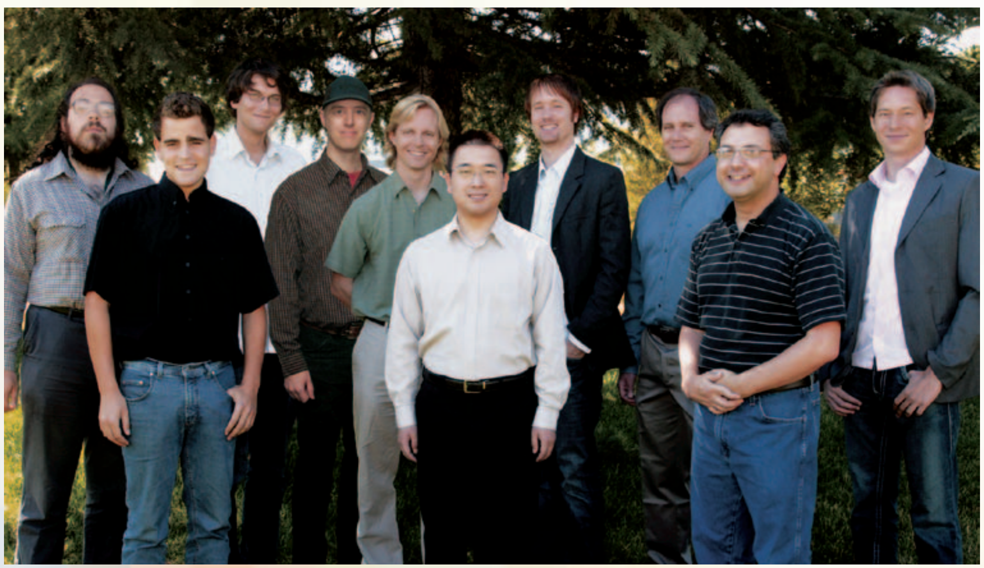

Development team for ROSE (from left): Peter Collingbourne, Martin Bauer (formerly of Livermore), Thomas Heller (formerly of Livermore), Robb Matzke, David Hamilton, Chunhua Liao, Andreas Saebjornsen, Daniel Quinlan, Jeffrey Keasler, and Thomas Panas.

Unique Industry–Government Collaboration

The Hyperion Project received a "Best HPC Collaboration Between Industry and Government" award from *HPCwire*, a news service dedicated to supercomputing. Funded by NNSA's ASC Program, Hyperion is an LLNL collaboration with industry partners to advance next-generation Linux high-performance computing (HPC) clusters. Collaborators include Dell, Inc.; Intel Corporation; Super Micro Computer, Inc.; QLogic Corporation; Cisco Systems, Inc.; Mellanox Technologies, Ltd.; DataDirect Networks, Inc.; LSI Corporation; Red Hat, Inc.; and Sun Microsystems, Inc.

Mark Seager, assistant department head for Advanced Technologies; Tom Tabor, publisher of *HPCWire*; and Matt Leininger, deputy director for Advanced Technologies, at SC09. Tabor presented LLNL and its 10 Hyperion Project collaborators with award certificates.

Top Executive Honors

Mark Seager, LLNL assistant department head for Advanced Technologies, was selected by *Federal Computer Week* as one of the "Federal 100" top executives from government, industry, and academia who had the greatest impact on government information systems in 2008. Seager was selected because of "the difference [he] made in the way agencies, companies and government officials develop, acquire, manage, and use information technology." The nomination was submitted by industry collaborators for Seager's leadership of the Hyperion Project. (See page 10.)

Mark Seager

Tom Slezak

Panayot Vassilevski

Three Patents for DNA Signatures

Three patents were issued to Tom Slezak, leader of Livermore's bioinformatics group, and several former LLNL employees for DNA signatures for the Class A biothreat agents *Francisella tularensis* (causative agent of tularemia), *Yersinia pestis* (causative agent of plague), and *Brucella* species (causative agent of brucellosis). The signatures were developed in 2000–2001 for the Biological Aerosol Sentry and Information System (BASIS) and were originally intended for the 2002 Winter Olympic Games in Salt Lake City. These are the latest DNA signature patents stemming from the early computational DNA signature work that led to LLNL's fully automated signature pipeline, KPATH, in 2002.

Fulbright Scholar

Computational mathematician Panayot Vassilevski was awarded a Fulbright Scholarship and served for six months as a senior lecturer in mathematics at the University of Sofia, Bulgaria. This prestigious international award was granted by the U.S. Department of State and offers lecturing/research awards in approximately 140 countries to college and university faculty and administrators, professionals from business and government, and artists, journalists, scientists, lawyers, and independent scholars. While at the University of Sofia, Vassilevski taught a graduate course based on his book *Multilevel Block Factorization Pre-conditioners*, which describes matrix-based analysis and algorithms for solving finite element equations.

Conservation Measures Add Up to Big Energy Savings

A two-year effort by the Computation Directorate to conserve energy by raising the temperature and managing air flow in the Terascale Simulation Facility (TSF) computer rooms is yielding big dividends—an estimated $2.4 million in energy savings annually so far. Called "Megawatts to PetaFLOPS," the program earned a 2009 DOE/NNSA Federal Energy Management Program Award. The savings came from a multifaceted approach, including raising the temperature in the computer rooms and the water used to cool some machines and finding efficiencies in the configuration of HPC systems.

TSF also received a Leadership in Energy and Environmental Design (LEED) gold level certification under the U.S. Green Building Council rating system. LEED is an internationally recognized certification system. It provides third-party verification that a building or community was designed and built using strategies aimed at improving performance in energy savings, water efficiency, carbon dioxide emissions reduction, and other factors.

"This is truly a noteworthy achievement for NNSA and LLNL that symbolizes our commitment to transforming the Cold War-era nuclear weapons complex into a modern, efficient nuclear security enterprise," said Brigadier General Garrett Harencak, NNSA principal assistant deputy administrator for military application, in a congratulatory communication.

Marriann Silveira and Anna Maria Bailey use a thermal imaging device to measure air temperature in the Terascale Simulation Facility's computer rooms.

Completed in late 2004, TSF is a 253,000-square-foot building that houses some of the world's fastest supercomputers, including Dawn (BlueGene/P), BlueGene/L, and Purple—ASC systems largely dedicated to stockpile stewardship. TSF represents an innovative design that emphasizes function over form.

RETURN TO TABLE OF CONTENTS

High-Performance Computing

MICHEL MCCOY,
MCCOY2@LLNL.GOV

More than ever before, the high-performance computing (HPC) landscape at Lawrence Livermore National Laboratory (LLNL) in 2009 has been a tapestry of competing themes, triumphs, and thrusts. This period has been less about acknowledging past accomplishments and more about building a future based on new formulas.

The most obvious accomplishment was the installation and rapid integration of Dawn, an IBM BlueGene/P system. Dawn is the initial delivery system of the Sequoia contract; Dawn will be employed both for code development in preparation for Sequoia and as a production resource in its own right. The 500-trillion-floating-operations-per-second (teraFLOPS) system was in pieces in January and February, but by March it was running calculations, culminating with a full-system laser–plasma calculation of a quad of beams propagating through a hohlraum for the National Ignition Facility (NIF). By fall, Dawn was moved to the classified network and made available to service the tri-laboratories (Lawrence Livermore, Los Alamos, and Sandia national laboratories). The Dawn integration was the most rapid of any system at Livermore Computing. Dawn is managed such that teams can test codes during regular business hours, but during nights and weekends it is used for batch calculations.

Dawn is a second-generation BlueGene system. The first-generation system, BlueGene/L, was the subject of national attention this year when it won the National Medal of Technology and Innovation. The research and development of BlueGene/L was mostly funded by the National Nuclear Security Administration (NNSA) at the behest of LLNL, and LLNL managed the contract with IBM. In a fortunate stroke of serendipity, Berni Alder, the LLNL scientist credited with inventing computational molecular dynamics, received the National Medal of Science the same year that BlueGene/L was recognized.

In 2009, Computation also procured and integrated a 100-teraFLOPS cluster, called Graph, which will act as a visualization engine for Dawn. Graph is already so reliable that some users prefer to work on it rather than the more established 100-teraFLOPS Purple system.

Computation is preparing for the 2012 arrival of Sequoia, the third-generation BlueGene system, which will be capable of performing 20 quadrillion floating operations per second (petaFLOPS). The preparatory work involves continuously reviewing the contract with IBM engineers and making complex trade offs about the system configuration. While this process has not been without challenges, the overall prognosis is that Sequoia will serve as a unique and powerful uncertainty quantification and science engine for the Advanced Simulation and Computing (ASC) Program.

In 2009, the Hyperion test bed won the *HPCwire* Editors Choice Award for Best HPC Collaboration Between Government and Industry. The Hyperion test bed will gain another partner in 2010, bringing the total companies on board to 12. LLNL will begin experimenting with flash memory in partnership with Fusion-io. Hyperion is, from LLNL's perspective, increasingly focused on testing technologies that will benefit the Sequoia system.

Looking forward, Computation is contemplating two complementary but decidedly different approaches to serving the ASC Program and staying relevant in the HPC community. The first approach concerns an exascale computing initiative, and the second concerns the development of an open campus.

In part due to the success of the second BlueGene research and development contract, funded both by NNSA and the Office of Science, Department of Energy (DOE) headquarters initiated a cross-DOE effort to make the case for a massive decadal investment in exascale technologies. LLNL is represented on the steering committee and is deeply involved in making the requirements and technology case for the multibillion-dollar exascale initiative. In addition, Argonne, Berkeley, and Lawrence Livermore national laboratories are seeking to form an institute focused on exascale computer science challenges, using key applications as drivers of innovation.

Computation sees the nature of leadership in simulation as drifting from the previously unassailable metric of peak FLOPS to more the role of partnering with companies to develop and disseminate simulation to enhance American competitiveness in the 21st century. To this end, Computation is working to establish a center in the Livermore Valley Open Campus and is addressing many related legal, business, and funding issues. If successful, we are confident that we will have developed an approach that will garner national attention.

In short, we are pleased with our many accomplishments in 2009, but our attention is as much focused on what lies five years down the road as it is on what is just ahead of or just behind us.

RETURN TO TABLE OF CONTENTS

New Dawn System Speeds the Way to the Sequoia Computing Era

Summary

Dawn is an IBM BlueGene/P computer system that was deployed at LLNL in 2009 as an initial delivery preview of the Sequoia 20-petaFLOPS system, which will be delivered in 2012. Initial delivery systems are a standard practice in major ASC computer procurements to help applications developers and users anticipate and prepare for a new system's expanded capabilities and features.

Dawn far surpasses the capabilities of the previous BlueGene systems at LLNL. Its new features include expanded memory, multicore processors, and operating system enhancements that allow for multiple threads of execution and dynamic loading of libraries. Most importantly, Dawn delivers expanded computing capability for the NNSA ASC Program. Scientists at IBM, and Lawrence Livermore, Sandia, and Los Alamos national laboratories have already run significant calculations on Dawn that delivered groundbreaking science, enhanced code performance, and some of the largest, highest resolution simulations ever performed in their respective scientific field.

Progress in 2009

The Dawn IBM BlueGene/P computer system consists of 36,864 four-way compute nodes, totaling 147,456 PPC450 cores, with 4 gigabytes (GB) of memory per node. The BlueGene microkernel operating system provides a stripped-down Linux environment to the compute nodes of the system, which minimizes system noise and allows applications to efficiently scale to the system's very large size.

Dawn has several new features, relative to the previous BlueGene systems, which are crucial to its role as precursor to the Sequoia system. Dawn's multicore processors and support for multiple threads of execution allow applications developers to explore new ways of expressing parallelism in their programs beyond the pure message-passing model. A hybrid style of programming that combines message-passing with OpenMP or other threading methods is considered essential to reach the full scale and capability of the anticipated million-plus-processor Sequoia system. Livermore Computing has initiated the Scalable Applications Preparation project to assist code teams with investigating hybrid methods, and Dawn is an essential tool in that effort. Another key feature of the Dawn architecture is its larger memory size per node, which allows the machine to be used by a number of applications that could not fit in the memory of the smaller BlueGene/L nodes. Dawn also supports dynamic loading of libraries, which benefits applications that employ Python as a programming language.

The 36-rack Dawn system was installed according to a very aggressive schedule. The first 18 racks were delivered on January 20, 2009, and acceptance of the full system was completed on March 26, 2009. LLNL and IBM devised an innovative plan to allow testing of the system as nodes were added, which greatly sped identification and resolution of issues and led to the successful deployment of Dawn in minimal time.

LLNL's ARES team demonstrated the expanded accessibility of Dawn to more sophisticated and demanding codes when they combined improved communications algorithms in their code with Dawn to complete a scaling study of a Zrad3D test case. The team achieved excellent scaling performance to the full machine size and demonstrated that threading is viable even at full scale on Dawn.

In one of its earliest calculations, Dawn delivered significant results for the pF3D code team in a calculation supporting NIF. The code team was able to model beam propagation of a 30° quadrant of the NIF point design, perform a whole-beam

Contact

SCOTT FUTRAL,
FUTRAL@LLNL.GOV

simulation with a high-fidelity physics model, and demonstrate acceptably low levels of backscatter for one particular laser configuration.

LLNL researchers applied innovative computer science techniques on the Dawn system, and their efforts resulted in award-winning calculations. The ddcMD team used new capabilities of their code and the Dawn computer to study plasma–beam interactions on a system of more than 400 million charged particles. This calculation provides insight into interactions that might occur in a fusion fast-ignition experiment. The ddcMD team's work was selected as a finalist for the 2009 Gordon Bell Prize. (See page 14.)

Dawn also serves as a magnet for external collaborations. Researchers from IBM and Lawrence Berkeley National Laboratory used computer time on Dawn to carry out a cortical simulation at unprecedented scale. The simulation had 900 million neurons and 9 trillion synapses and used probabilistic connectivity. These runs exceeded the scale of a cat brain cortex, and the resulting paper was selected as the 2009 Gordon Bell Prize winner.

As a tri-laboratory resource for the ASC Program, Dawn has delivered important results for Los Alamos and Sandia national laboratories. Sandia researchers used the PDLAMMPS code on 65,000 cores to simulate glass dynamic brittle fracture and, in doing so, performed the largest peridynamic simulations in history. Sandia researchers also ran CCSM-HOMME, an experimental configuration of the Community Climate System Model, and achieved higher resolution simulations at a record-setting rate. Researchers from Los Alamos conducted a Rayleigh–Taylor instability calculation that is thought to be the largest fully resolved simulation of its kind ever performed at a high-density ratio using more than 77 billion grid points.

The 8th-century BCE Greek poet Hesiod said, “Dawn speeds a man on his journey, and speeds him too in his work.” This quote appropriately summarizes the intent of the Dawn system. By developing innovative technologies and algorithms in anticipation of the Sequoia system, computer scientists are using Dawn to speed the way to the future and deliver important scientific results today.

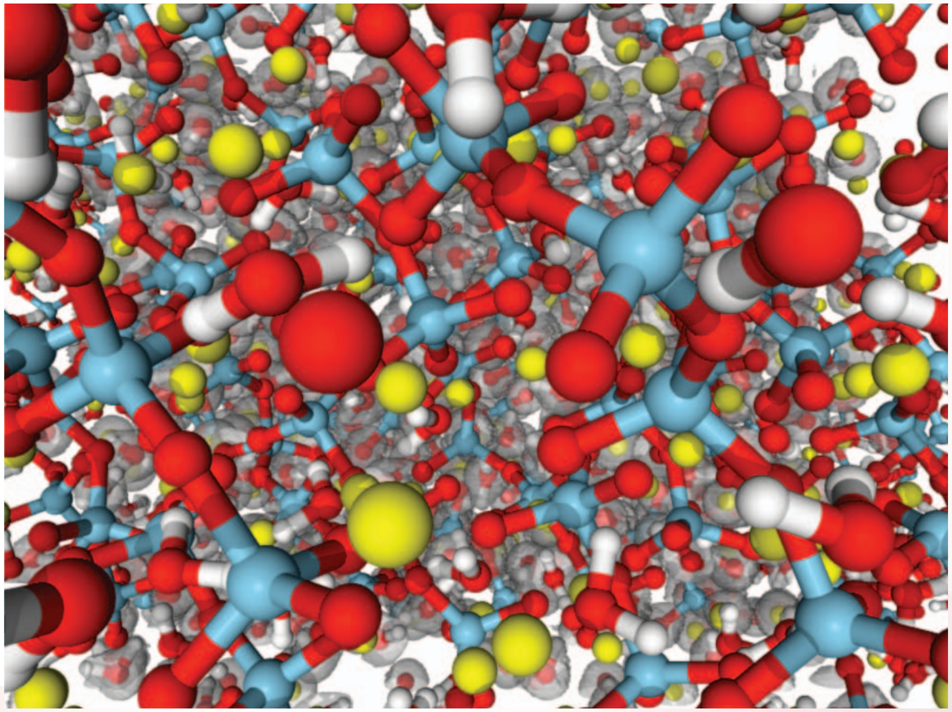

LLNL scientists used Dawn to complete first-principles molecular dynamics simulations of a large-scale atomistic model for cement using the Qbox code. Early results showed efficient scaling up to 16K nodes, and preliminary timing of larger model systems exhibited good scaling on larger partition sizes.

Hyperion Represents a New Paradigm for Advanced Computing Technologies

Contact

Matt Leininger,
matt@llnl.gov

hyperionproject.llnl.gov

Summary

One of the challenges associated with developing, managing, and utilizing HPC systems is a lengthy and costly testing phase that must be completed to ensure the new hardware and system software works well at scale. LLNL and 10 industry partners—Dell, Intel, Super Micro, QLogic, Cisco Systems, Mellanox, DataDirect Networks, LSI, Red Hat, and Sun Microsystems—have developed an innovative solution to this problem. Hyperion is a large-scale Linux cluster test environment used to speed the development of HPC clusters by making a cluster available solely for developing and testing infrastructural components. It is also used to evaluate system hardware and software critical to maintaining the aging U.S. nuclear weapons stockpile. The Hyperion Project was recognized with a 2009 Federal 100 award and an *HPCwire* award for Best HPC Collaboration Between Government and Industry.

Progress in 2009

Hyperion, too large for any one organization to deploy, represents a unique approach for industry-leading collaborations and exemplifies LLNL's leadership in HPC. Hyperion is a dedicated development and test environment for advanced Linux cluster technologies, an evaluation test bed for new hardware and software technologies, and an input and output (I/O) test bed to prepare for the future ASC Sequoia platform.

The second phase of Hyperion was installed in March 2009, and it doubled the system size to 1,152 nodes. The first half of the cluster uses two-socket quad-core Intel Harpertown processors, 8 GB of memory per node, and a double data rate (DDR) InfiniBand interconnect, while the second half uses quad-core Intel Nehalem processors and 12 GB of memory per node. The system expansion allows for more extensive testing of the performance and scalability of critical HPC technologies.

Hyperion is being used to test the Lustre parallel file system, the Red Hat-based ASC Tripod Operating System Software, and the OpenFabrics high-performance networking software for InfiniBand and low-latency Ethernet. Several high-performance storage technologies from DataDirect Networks, LSI, Sun Microsystems, Terascala, and Xiotech are undergoing evaluation, as are two separate storage area networks (SANs)—one based on InfiniBand and another based on low-latency 10-gigabit Ethernet.

In 2009, Hyperion was moved to the LLNL collaboration ("Green") network to allow a more diverse group of users access to the system. In addition, Memorandums of Understanding were signed between LLNL and the 10 industry partners.

Next year, Hyperion storage and SAN environments will be expanded to evaluate emerging technologies. Disruptive storage technologies involving high-performance solid-state storage and storage server virtualization will be used to design data-oriented architectures. These architectures are expected to benefit scientific application checkpointing and data analysis. Solid-state storage will provide high I/O operations per second and high bandwidth to I/O-dominated applications. Virtualization will allow the data analysis algorithms to run on the storage server node closest to the data, thereby reducing network traffic and latency. The Hyperion SAN will be upgraded to include quad data rate InfiniBand with 40-gigabit-per-second bandwidth, adaptive routing, and network congestion control features. The successful storage and SAN technologies will be integrated into the ASC Sequoia simulation environment.

The Hyperion cluster architecture includes 1,152 nodes using an InfiniBand DDR interconnect. Hyperion connects to two SANs—one based on Ethernet and the other on InfiniBand.

Return to Table of Contents

Checkpointing Tool Improves File Bandwidth for Application Fault-Tolerance

Contact

Adam Moody,
moody20@llnl.gov

Summary

The Scalable Checkpoint/Restart (SCR) library delivers high bandwidth and substantially reduces the cost to checkpoint large-scale simulations. SCR allows simulations to run more efficiently, recover more work in the event of a system failure, and reduce the load on critical shared resources, such as the parallel file system and the network infrastructure. Initially conceived and developed by Livermore Computing to improve service to a particular simulation, SCR was expanded in 2009 and is now offered to all users of Livermore Computing's Linux clusters.

Progress in 2009

Today's supercomputers consist of hundreds of thousands of hardware components. Occasionally an underlying hardware component fails, which interrupts the simulations running on the machine. To be able to progress in the presence of hardware failures, simulations periodically save their state in a checkpoint. Then, when a failure occurs, the simulation restarts from its most recently saved checkpoint. Checkpoints are traditionally stored in a parallel file system, which can be costly.

SCR's solution is derived from two key observations. First, a simulation only needs its most recent checkpoint. As soon as the next checkpoint is written, the previous checkpoint can be discarded. Second, a typical failure disables a small portion of the system but otherwise leaves most of the system intact.

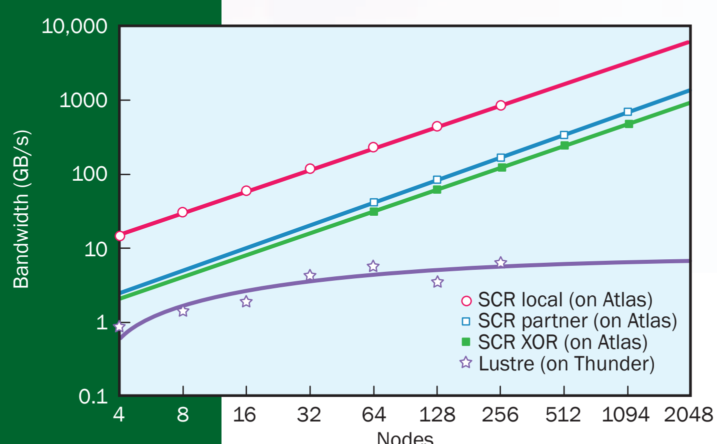

Aggregate file-write bandwidth via the parallel file system (Lustre) and via SCR's three cache redundancy schemes.

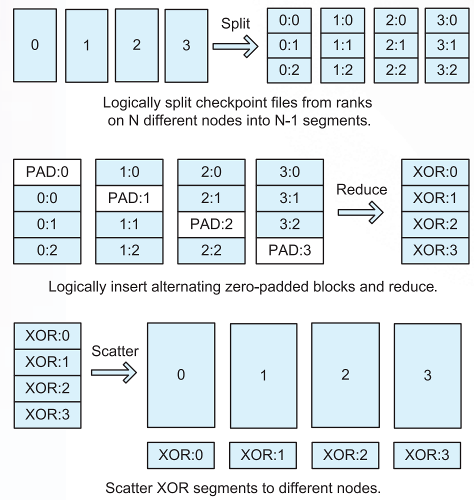

XOR redundancy scheme applied to files in the checkpoint cache.

SCR is designed to cache checkpoints in storage that is local to the system rather than in a parallel file system. SCR caches only the most recent checkpoint; the previous checkpoint is discarded when a new checkpoint is saved. SCR also applies a redundancy scheme to the cache, making it possible to recover checkpoints even after a failure disables a small portion of the system. This design allows each checkpoint to be written to the cache and discarded unless a failure occurs, at which point SCR recovers the most recent checkpoint from the cache and copies it to the parallel file system. By writing to the cache instead of the parallel file system, SCR delivers exceptional bandwidth and perfect scalability on LLNL supercomputers.

In 2009, the SCR application programming interface and implementation were expanded to support a wider range of simulations. A new redundancy scheme, XOR, was implemented to reduce the cache storage requirements, enabling more simulations to utilize SCR with a random access memory disk. Also, to allow simulations with large memory requirements to use SCR, checkpoints can now by cached in hard drives or solid-state drives on systems where such hardware is available. SCR was open-sourced this year to enable collaborations with other computer centers.

SCR offers several benefits. In one case, customers running large-scale simulations saw a decrease in the time required for checkpointing from 20 minutes to 15 seconds, an improvement of 27% in machine efficiency during failure-free periods, a decrease in the average time lost due to a failure from 85 minutes to 12 minutes, and a seven-fold reduction in data written to the parallel file system.

RETURN TO TABLE OF CONTENTS

Exascale Computing Technologies Improve System Productivity

Summary

Current trends in computer architecture are leading toward large-scale systems with more than 1.5 million processors, such as the Sequoia supercomputing system under development at LLNL. In a slightly longer time frame, exascale systems that process over 10^{18} FLOPS (or 1 exaFLOPS) are expected to require between 500 million-way to 4 billion-way parallelism. Computing systems will reach this extraordinary scale by using multicore chips, resulting in less memory and less memory bandwidth per core. Exascale systems will also include novel hardware features, such as thread-level speculation and transactional memory. The Exascale Computing Technologies (ExaCT) Strategic Initiative is working to address the challenges arising from exascale systems and provide the innovations these systems will demand. Researchers are designing scalable mathematics algorithms and application-level fault tolerance techniques that can be adapted to emerging architecture trends. In addition, the innovative code correctness and performance methodologies will automate critical aspects of application development. The new research will significantly increase the scientific output obtained from LLNL's large-scale computing resources by improving application scientist productivity and system utilization.

Progress in 2009

The ExaCT Strategic Initiative encompasses several interrelated thrust areas in computer science and mathematics: code correctness and performance methodologies, application-level fault tolerance techniques, and scalable mathematics algorithms for multicore systems. Innovations in code correctness and performance methodologies will exploit the fault tolerance techniques. Similarly, the results attained in implementing the code correctness and performance technologies will guide development work on novel mathematics algorithms, which will build on the fault tolerance techniques.

ExaCT researchers are exploring data layouts to substantially improve algorithm performance. They also are developing automated mechanisms to extend these results to general LLNL simulations. The team is designing multigrid solvers with fine-grain parallelism and will pay careful attention to their convergence properties to support building fault tolerance into the algorithms. In 2009, significant progress was made in developing scalable algorithms that target multicore nodes. In the past, application scientists chose between an old, high-complexity operator, which offers good convergence properties, and a low-complexity operator that improves memory usage but requires many more iterations to achieve convergence. The new algebraic multigrid interpolation operator delivers improved convergence properties for very-low complexity coarsening schemes. Application-level fault tolerance methodologies developed by the ExaCT team can also be adapted to improve general applications. For example, the 2009 results with pF3D show that in-memory checkpointing, which records the state of an in-progress simulation, increases throughput by 30%. However, to apply these results to future systems, the team must overcome problems associated with having less memory per processor.

In addition, ExaCT researchers developed a three-pronged approach for their code correctness methodologies. First, lightweight mechanisms will efficiently gather data from application processes at the system's full scale and identify behavior equivalence classes, sets of processes with similar behavior. By narrowing the search for root-cause analysis to a few nodes, researchers can reduce problems to a scale suitable for traditional interactive techniques. Second, the team is extending scalable techniques to track and verify user expectations of algorithmic properties, such as convergence rates. Third, they are developing several scalable techniques to automatically identify the source of a defect and the type of code or algorithm causing the problem. Overall, the team's

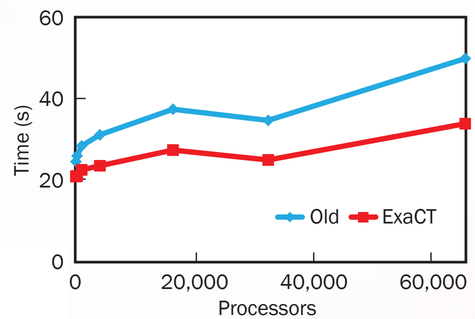

Processors	High	Old	ExaCT
512	13	49	17
4,096	16	75	21
8,000	18	85	21
15,625	26	95	22
32,768	25	107	23
Operational complexity	3.17	1.22	1.37

Scaling of interpolation operations (rotated anisotropy, 0.01, 60°, 500 × 500 cores per processor on the unclassified BlueGene/L system). The table lists the computing operations required for convergence on the three types of operators. The figure compares the performance offered by the new operator to that of the old operator and shows that this performance advantage increases with the number of processors.

CONTACT

BRONIS DE SUPINSKI,
BRONIS@LLNL.GOV

approach will reduce the time scientists spend analyzing code to determine which module, whether in an application or system software (e.g., the Message Passing Interface library), contains the root cause of a problem that arises at scale.

In 2009, the ExaCT team also began work on a debugging framework. This technology will collect data from a running application and store it in a persistent repository. Information in the repository can then be used to determine the application's state and whether its behavior is correct. The team is investigating what data sources best support equivalence-class identification and how to extract such data from an application without compromising scalability, incurring significant overhead costs, or introducing load imbalance.

ExaCT researchers are adapting a similar infrastructure to address application load balance issues. In 2009, they developed a scalable mechanism to measure how an application varies over time in key dynamic execution regions. They will extend this mechanism not only to measure application loads but also to identify imbalances automatically and to calculate the response actions and time needed to correct them.

An important part of the ExaCT Strategic Initiative is to test research results on LLNL's mission-critical applications. To this end, the team is working with production-scale applications developed by scientists in LLNL's Weapons and Complex Integration and NIF and Photon Science principal directorates to ensure that the results are relevant to the missions of LLNL and DOE. Ultimately, in anticipating the dramatic changes in future large-scale computing systems, LLNL will deliver the innovative advances that application teams need to successfully exploit these resources.

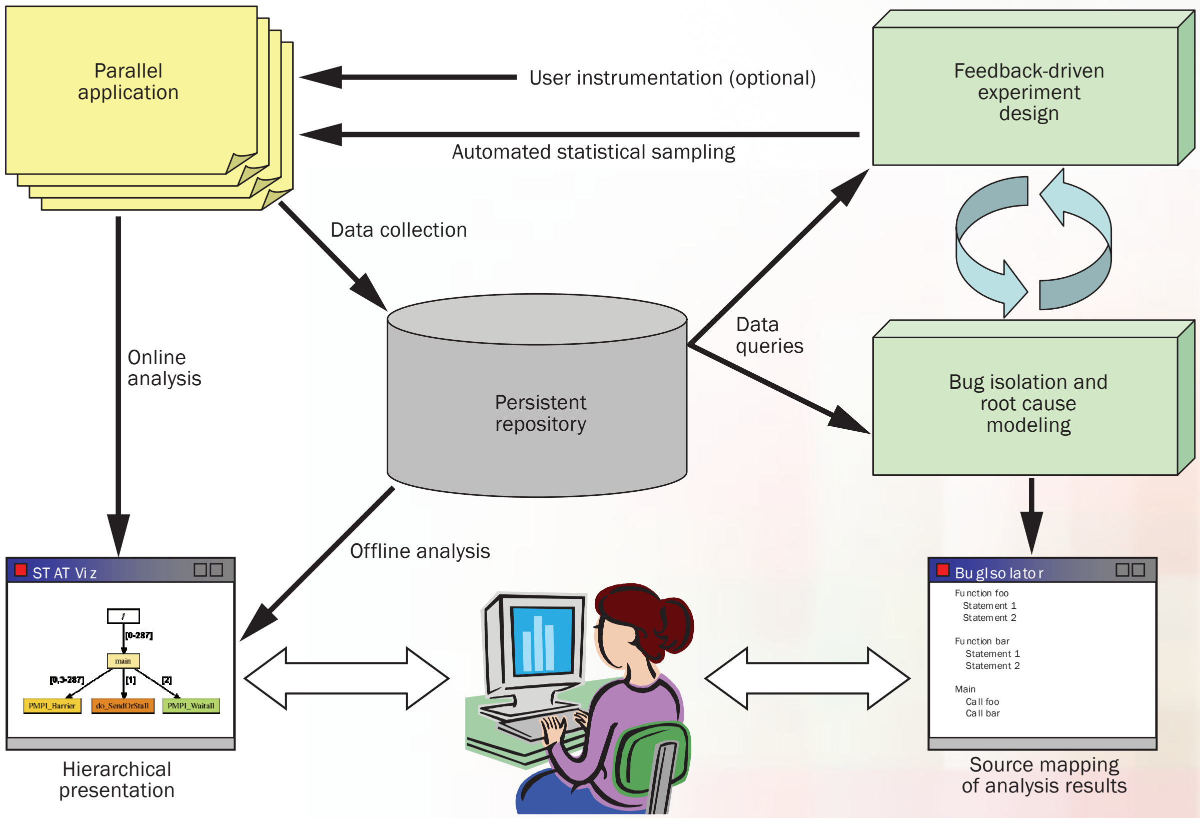

Code correctness and performance analysis architecture designed as part of the ExaCT Strategic Initiative.

RETURN TO TABLE OF CONTENTS

CONTACT

FRED STREITZ,
STREITZ1@LLNL.GOV

ISCR.LLNL.GOV

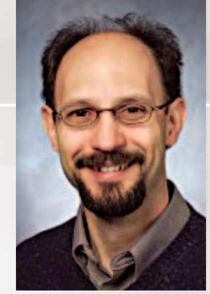

ISCR Delivers Plasma Properties at Unprecedented Scale

Summary

A team of scientists from the Institute for Scientific Computing Research (ISCR) demonstrated an innovative parallelization strategy by simulating hot, dense plasma at unprecedented resolution and scale. Their strategy helped solve one of the most difficult scaling problems in numerical simulation: the efficient parallel calculation of long-range interactions. Long-range forces (such as electrostatic or gravitational) are relevant to a variety of modeling scenarios, but prior efforts to develop a fully scalable solution failed on processor counts beyond a few thousand. The team's ability to simulate 10 billion particles across more than 290,000 processers delivered game-changing capability to LLNL's Weapons and Complex Integration and NIF and Photon Science principal directorates. The team was recognized as a finalist for the 2009 Gordon Bell Prize.

Progress in 2009

The temperature and pressure of a heated and compressed plasma as it approaches fusion conditions depend sensitively on the relative rates of energy exchange between the target electrons and the nuclei. A complete simulation of the energy production and flow in a plasma fusion problem would be invaluable, but it is not yet possible due to the high complexity of the different physical processes that need to be included. The team initially focused on calculating the long-range (1/r) electrostatic, or Coulomb, interaction, as this piece is physically the most important and computationally the most demanding aspect of large-scale plasma molecular dynamics simulations. The key insight was to pay particular attention to the scaling characteristics of the individual components in the simulation. Using a single strategy for parallelization was limiting the overall scalability of the calculation because the entire code was held back by the least scalable component. By decomposing the calculation across the computer in a very heterogeneous fashion (i.e., assigning different parts of the calculation to different processors), the team exploited the ability to calculate real-space and Fourier-space force terms concurrently, which allowed the overall calculation to scale effectively across hundreds of thousands of processors.

The team's innovative approach of breaking up the machine to run the calculation has far-reaching implications for scientific computing and will likely affect the way future codes are developed. The flexibility of their approach allows more complicated models to be developed, and the technique can be applied to other applications on both current and next-generation machines. The team plans to include additional, shorter-range physical processes (e.g., radiation, recombination, and fusion) in the code, with a goal of delivering a comprehensive simulation tool for computing correlations and transport properties in burning plasma.

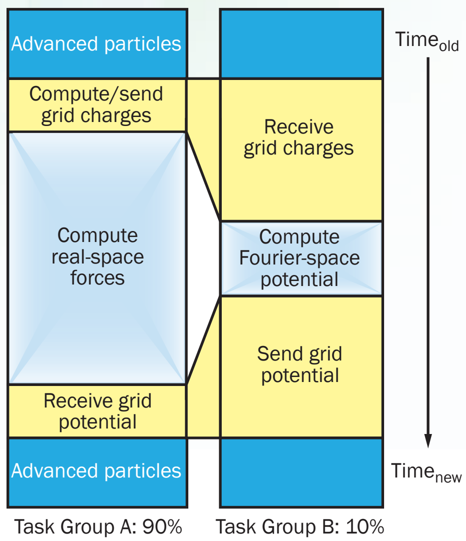

Schematic of the time and work flow used to scale the calculation of long-range forces across a massively parallel supercomputer. Note that although Task Group B spends most of the time in communication, these processors represent only a fraction of the machine. The majority of the processors (Task Group A) spend most of their time in computation.

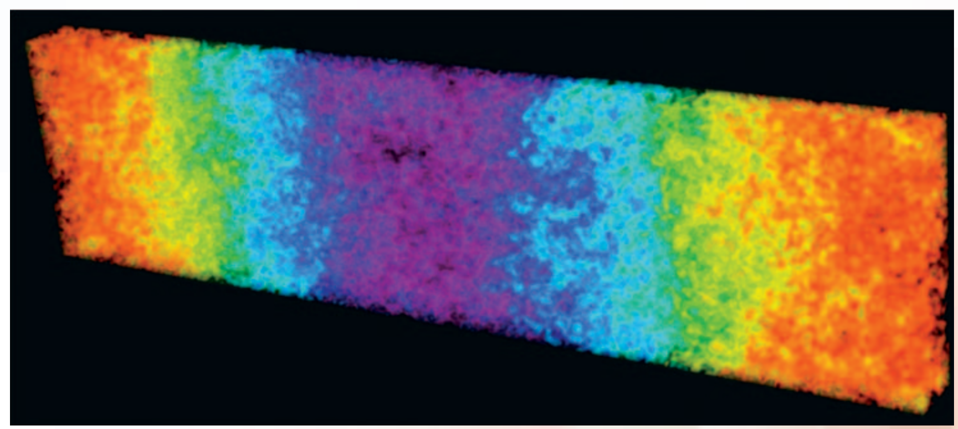

Electron temperature in an argon-doped deuterium–tritium plasma (initially in equilibrium at 500 electronvolts and a density of 1025 per cubic centimeter) being heated by a proton beam. The 400-million-particle simulation was performed on 147,456 processors of the ASC Dawn computer at LLNL.

RETURN TO TABLE OF CONTENTS

Synthetic WorkLoad Provides Enhanced Acceptance Testing

CONTACT

SHEILA FAULKNER,
SAF@LLNL.GOV

SUMMARY

Synthetic WorkLoad (SWL) is a collection of functional tests, performance benchmarks, and user applications developed by Livermore Computing for acceptance testing of new computer systems. This testing suite simulates a typical production workload as part of the process for evaluating a computer system before it is placed in production mode. SWL also includes a framework for managing the tests, generating and executing sample workloads, and reporting the results. In 2009, Livermore Computing extended the SWL suite to support other scenarios when hardware and software must be validated before a system returns to production mode.

PROGRESS IN 2009

A typical SWL sequence begins with a set of functional and performance tests to verify that a new system can execute a series of simple functions to the level required by the procurement contract. If the functional test phase is successful, SWL generates a workload that the new system runs for three to five days. During this stability-phase testing, the system must maintain certain utilization levels and operate without any major failures.

SWL has played an important role in the system checkout and acceptance testing for most of the large clusters available at Livermore Computing, including the Peloton and Tri-Laboratory Capacity Cluster (TLCC) procurements. The TLCC procurement is noteworthy because it included systems at Lawrence Livermore, Los Alamos, and Sandia national laboratories. Personnel at the three laboratories used SWL to check those systems before shipment and after final installation.

In 2009, SWL was extended to include a series of the Tier 1 benchmarks for Sequoia, the 20-petaFLOPS IBM system due in 2012. Many SWL management scripts had to be upgraded to support the multiple system configurations and compilers on Dawn, an IBM BlueGene/P system, and Graph, a large Linux cluster designed for visualizing the simulations run on Dawn.

To improve the coverage for the 500-teraFLOPS Dawn computer—the precursor system for Sequoia—developers reconfigured several tests to run in multithreading (OpenMP) mode in addition to the usual parallel mode required by MPI, the message-passing interface. This upgrade allowed developers to examine and characterize the multithreading capabilities that are part of the BlueGene/P architecture.

SWL's job creation and execution scripts were also modified so SWL could examine the 24 cores in every Graph node. In addition to this expanded number of processing elements, each node contains 128 GB of memory. To more fully examine the increased memory, developers extended the configurations on many SWL tests.

Another 2009 accomplishment was expanding the SWL suite to cover other testing needs. Developers modified several functional tests to run in single-node configurations so they could be used to examine recently repaired hardware and other troubleshooting scenarios. SWL was also adapted for regression testing after a major hardware or software upgrade as well as for prerelease testing on new versions of CHAOS, Livermore Computing's clustered high-availability operating system. The CHAOS testing is particularly important because this application is part of the Tripod Operating System Software released on the TLCC systems.

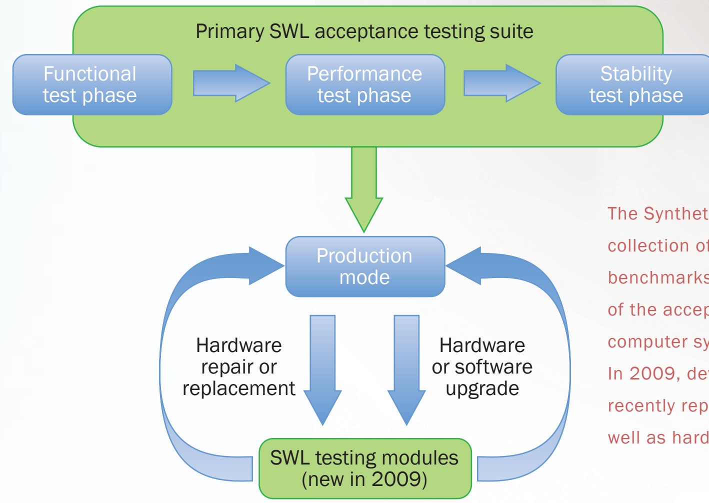

The Synthetic WorkLoad (SWL) suite runs a collection of functional tests, performance benchmarks, and user applications as part of the acceptance testing before a new computer system is placed in production. In 2009, developers extended SWL to test recently repaired or replaced hardware as well as hardware and software upgrades.

RETURN TO TABLE OF CONTENTS

SLURM Advances Resource Management

Contact

MOE JETTE,
JETTE1@LLNL.GOV

COMPUTING.LLNL.GOV/LINUX/SLURM

Summary

At a growing number of computer centers around the world, high-performance computing is done on massively parallel systems running the Simple Linux Utility for Resource Management (SLURM). As a highly scalable resource manager, SLURM executes parallel applications, manages the queues of pending work, allocates resources to those jobs, and provides a framework for starting and executing each job. SLURM began in 2002 as an open-source project at LLNL. Since then, Hewlett-Packard, Bull, and many other computing industry and academic partners have contributed to the utility's success. In 2009, major enhancements were made to SLURM's job-scheduling capabilities. An innovative logic process optimizes resource allocations with respect to a computer's network topology, and an integrated mechanism automatically powers down idle resources to save energy and reduce operating costs.

Progress in 2009

SLURM has a modular design, offering users dozens of options to support various architectures and enhance the utility's flexibility. New modules in version 2.0, which was released in May 2009, extend SLURM's capabilities, advancing it from a simple resource manager (with first-in, first-out queues) into a sophisticated job scheduler. This upgrade moves important functionality, previously provided by commercial applications, directly into SLURM.

For example, an assortment of resource limits, such as maximum job time or size, can be imposed on individual users or hierarchical bank accounts. Pending jobs can be prioritized according to such criteria as age, size, or proximity of users or banks to their allocation targets. A hierarchical fair-share algorithm boosts the priority of pending jobs for users who are underserved relative to their allocation. In the default mode, SLURM assigns idle resources to pending work as running jobs complete, but it can also be configured so that high-priority simulations will preempt lower-priority work for faster service. Preempted jobs can be cancelled, suspended, requeued, or checkpointed, and a reservation capability ensures that resources are available when needed for critical applications or system maintenance.

Another new technique considers a computer's network topology when allocating resources to minimize the communication distance between a job's components, thus improving performance. For example, a space-filling Hilbert curve maps the resources required for a specific topology, such as a three-dimensional torus network, into a one-dimensional space. SLURM's native best-fit algorithm then operates on that defined space to assign resources with a high degree of locality.

In 2009, SLURM was also enhanced to completely power down idle resources and power them back up on demand, thus minimizing the energy costs associated with high-performance computing. This functionality is useful not only when a system has a light workload but also when it is accumulating resources for large, time-consuming jobs.

Ongoing SLURM development work includes porting to new computer architectures such as Sequoia, the 20-petaFLOPS IBM system due in 2012. Another major effort will enable SLURM commands to operate between clusters, permitting a user to check the status of a job or submit new jobs between clusters. Improvements such as these ensure that SLURM continues to provide stellar resource management capabilities for high-performance computing.

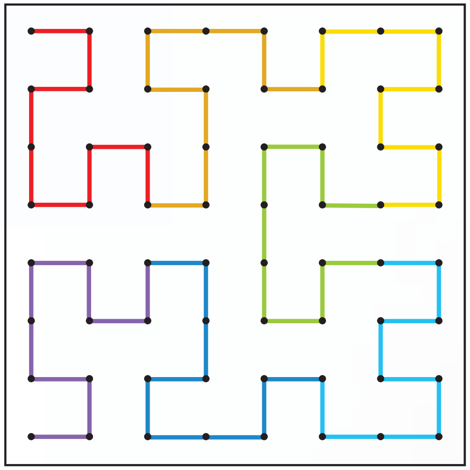

This Hilbert space-filling curve assigned an order to a set of 8 × 8 nodes that are collectively "near" one another in space. SLURM then assigns resources to a job's various components based on the resulting node order.

RETURN TO TABLE OF CONTENTS

Research and Advanced Development

LORI DIACHIN,
DIACHIN2@LLNL.GOV

Research and development in the Computation Directorate is focused on preparing Lawrence Livermore National Laboratory (LLNL) for a new era in computer and computational science. As we look toward the deployment of an exascale system in the next decade, it is clear that modeling and simulation activities at LLNL are facing a fundamental paradigm shift. We anticipate that the entire computing ecosystem, from architecture design to programming models, mathematical libraries, and applications, will look considerably different in a decade's time than it does today. Computer architectures will have many millions of processing cores and deep memory hierarchies, and they will require software that is resilient to faults, adaptable, and sensitive to power consumption. At the same time, the explosion of data associated with numerical simulations, experimental scientific facilities, and information sciences is driving an equally compelling research agenda in novel approaches to knowledge discovery. To meet these challenges, Computation Directorate research is focused on developing innovative computer science approaches and new mathematical algorithms and encapsulating the results of that research into software libraries and tools that can be used by domain scientists. If successful, we will enable better science across all Laboratory missions by delivering the foundations necessary to pursue computational and data science applications orders of magnitude greater in capability and capacity, thereby achieving results with significantly greater accuracy and fidelity to reality.

To accomplish our research objectives, the Computation Directorate carries out a program that spans many disciplines, including applied mathematics, computer science, and data science. By engaging in strategic planning in these areas, we strive to anticipate the needs of the Laboratory programs. In 2009, we were heavily engaged in the joint Department of Energy's (DOE) National Nuclear Security Administration (NNSA) and DOE Office of Science planning activities for exascale computing by providing expertise and leadership in a series of workshops and reports that will inform DOE investment decisions. We also made significant contributions to LLNL's 100-day strategic planning activity, which is guiding internal Laboratory Directed Research and Development (LDRD) investments in key application thrust areas and cross-cutting science pillars, such as high-performance computing. These and similar activities ensure that our research efforts benefit from expert analysis by the broader scientific community and meet the most critical needs of LLNL programs.

We pursue innovative long-term research in mathematics and computer science to meet programmatic needs through a mix of internal support (via the LDRD Program) and external support (e.g., DOE's Office of Advanced Scientific Computing Research). For example, Computation researchers have been leading or participating in LDRD strategic initiatives to develop algorithms and infrastructure for computer network defense, create new mathematical approaches for uncertainty quantification and error estimation for Laboratory simulations, and develop novel computer science tools that will allow us to assess performance, debug codes, and resume from faults on next-generation computing architectures. Many exploratory research LDRD projects are also under way, and some are highlighted in this report: evaluating novel architectures for storage intensive computing; developing a predictive approach that uses multiscale modeling, bioinformatics analysis, and Laboratory experiments to uncover the basic principles of viral evolution; creating efficient, high-fidelity continuum algorithms for the Vlasov-Maxwell systems that are critical to facilitating routine laser-plasma interaction simulations; and using parallel discrete event simulation technologies to develop a modeling capability for the space surveillance network. These examples help demonstrate the breadth and depth of the Computation research portfolio.

Finally, a critical aspect to Computation research is our focus on delivering results to LLNL missions through internal programmatic support (e.g., NNSA's Advanced Simulation and Computing Program) and to the broader scientific computing community (e.g., with funding from the Office of Science's Scientific Discovery through Advanced Computing Program). For example, the SAMRAI software for structured adaptive mesh refinement, which has origins as a research project under LDRD auspices, is now a foundational component in the development of the next-generation ARES code. Moreover, Computation software tools for visualization, language interoperability, solution of linear systems, and performance improvement, among others, are released as open source and downloaded thousands of times a year. Such tools often receive external recognition; in 2009, the ROSE compiler infrastructure for performance enhancements and software analysis received an R&D 100 award.

As we move into what promises to be an exciting new decade for high-performance computing, Computation's research portfolio will continue to be driven by the cycle of anticipate, innovate, and deliver.

RETURN TO TABLE OF CONTENTS

Uncertainty Quantification Research Helps Refine Predictive Simulation

Contact

Xabier Garaizar,
garaizar1@llnl.gov

Summary

Uncertainty quantification (UQ) is an emerging branch in computational science designed to characterize and ultimately reduce the "uncertainty," or error, in predictive simulations. Simulating the effects of a possible future scenario, say, an increase in global temperatures over a 100-year timeframe or the aging of materials in a nuclear weapon, involves complex numerical calculations run on advanced computing systems. By quantifying the uncertainty in a given simulation, scientists and decision makers can better analyze predicted scenarios, much as the calculated margin of error helps frame the accuracy of experimental data. As part of the UQ process, scientists must better understand how discretization errors propagate and ultimately affect the predicted results. In 2009, Computation scientists focused on advancing three areas of UQ research: adaptive sampling, code convergence, and visualization.

Progress in 2009

A common UQ approach is to run ensembles of simulations that span the variation expected for the parameters most likely to dominate the uncertainties in a model. In 2009, Computation scientist Charles Tong investigated the use of adaptive sampling techniques to improve the robustness and efficiency of UQ and sensitivity analysis methods. Tong developed and assessed variable selection methods to identify relevant parameters that significantly influence model output uncertainties. In addition, he applied adaptive sampling to the construction of response surfaces, which approximate the outcomes of numerical simulations at different uncertain parameter values. He also used sampling refinement techniques to more efficiently quantify the sensitivities of the relevant uncertain parameters. These new methods are now included in PSUADE, an LLNL software package that assists scientists in identifying the most important parameters in a complex model and in studying how experimental data can help reduce the model uncertainties.

In another 2009 project, Computation scientists Carol Woodward and Jeff Hittinger investigated the code convergence behavior of several multiphysics simulations used by weapons researchers at LLNL. Code verification relies on error estimation techniques based on sequential reduction of the discretization parameters in a simulation so that numerical solutions will converge toward a unique and stable solution. Unfortunately, nonlinear couplings and competitions between the different models in a multiphysics code create ambiguities in this behavior. Woodward and Hittinger are developing better procedures to address the complications and subtleties of multiphysics code convergence. Their work builds on previous code and calculation verification studies completed by Computation scientists in support of the tri-laboratory Advanced Simulation and Computing (ASC) Verification and Validation Level 2 Milestone.

The UQ efforts planned for 2010 include a visualization and data-analysis project to overcome the difficulties caused by the size of simulation ensembles and the high dimensionality of the data involved. In addition, Computation researchers are participating in a multidisciplinary project to develop methods for reducing the dimension of the uncertainty parameter space in LLNL's multiphysics codes. This project will also investigate how to write future codes so that simulations provide not only an output value but also error bounds and distributions. In all of these projects, Computation will continue to work on improving the fidelity of LLNL's simulations and on better understanding how uncertainty affects the predicted results.

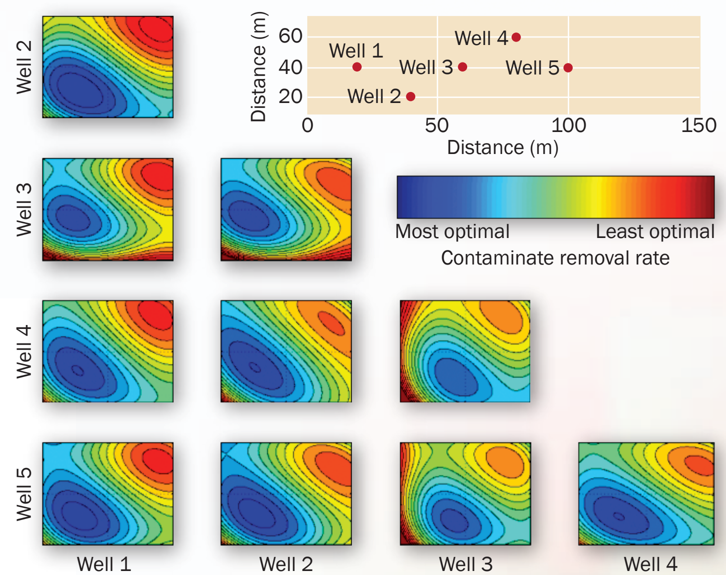

Uncertainty quantification methodologies, applied to a PSUADE calculation, show the best rate at which to pump contaminants from five wells while minimizing cost. Dark blue indicates the region with the optimum pumping rate at the lowest cost for a given well.

RETURN TO TABLE OF CONTENTS

Research Improves the Performance of Data-Intensive Applications

CONTACT

MAYA GOKHALE,
GOKHALE2@LLNL.GOV

SUMMARY

Across the Laboratory, and in the national security and scientific communities at large, analysts and researchers are searching for techniques, tools, and computing architectures to manage and analyze large data sets. Data-intensive problems are particularly common in text processing, sensor-related activities, and scientific data analysis. For applications that require frequent access to large data streams and stores, the traditional technology, which separates compute from storage resources, is inadequate. The Storage Intensive Supercomputing (SISC) Project improves the computation of data-intensive problems in national security and basic science by optimizing advances in architectures, systems software, and algorithms. The team's research into storage-coupled computing architectures and software addresses applications that cannot run on current systems. The team has demonstrated more than an order-of-magnitude improvement in performance over current systems.

PROGRESS IN 2009

In 2009, the SISC project defined a high-performance, deployable, data/compute-intensive node architecture that enhances the microprocessor with a coprocessor-accelerator as well as solid-state persistent memory. The team researched algorithm acceleration using many-core processors, such as the Ambric 336-core device and the Tilera 64-core system on a chip, and demonstrated an increase in speed of 20 to 60 times that of standard microprocessors. With multithreaded algorithms, solid-state flash memory can be accessed 2 to 3 times faster than rotating hard drives.

The SISC data/compute node architecture has been adopted by an LLNL advanced technology demonstration project involving an improvised explosive device detection system in a vehicle. The mobile compute platform, which is housed in a Chevy Suburban, uses a low-power microprocessor with an integrated graphics processing unit (GPU) and solid-state drive storage. It is capable of processing data in real time from an array of sensors, including radar, Light Detection and Ranging, and cameras.

In the area of systems software, SISC researchers are investigating a metadata-rich file system. In current data management systems, application developers store data sets in file systems and enter user-defined metadata describing data characteristics into a database, making it necessary to coordinate and synchronize disparate file and storage systems. In the new searchable metadata-rich file system, users attach arbitrary metadata to files and link files according to semantic relationships. For example, a Sloan Digital Sky Survey image can be tagged with the names and locations of brown dwarves found in the image. A user can query the file system for images tagged as "brown dwarf," creating a virtual directory linking all such images, and then initiate a processing chain to analyze the images.

The SISC team also upgraded the 140-node Fenix cluster in 2009. Originally created by the team in 2007, Fenix was the first LLNL data cluster to colocate storage with compute resources and has since gained popularity among scientists seeking to research storage-intensive software and applications. In addition to being continuously used for algorithms and systems research by scientists in the Center for Applied Scientific Computing, the newly dubbed Tuson data cluster is being used by SLAC National Accelerator Laboratory to prototype next-generation scientific data management software.

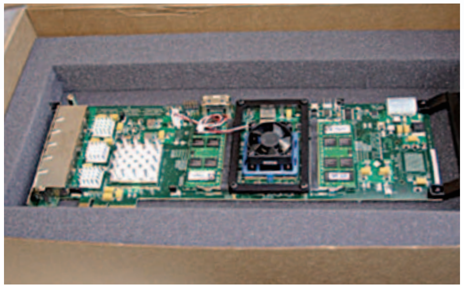

SISC algorithms utilize Tilera many-core processing boards.

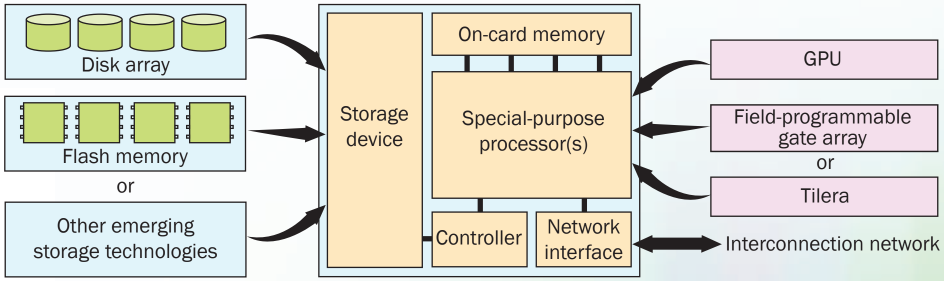

The SISC data/compute node architecture.

RETURN TO TABLE OF CONTENTS

New Analytical Paradigms Transform Cyber Security Research

Summary

The Supercomputing Enabled Transformational Analytics Capability (SETAC) Project is a three-year LLNL-funded Strategic Initiative to create fundamentally new analytical approaches for situational awareness applied to cyber security. SETAC began in November 2008 and will be complete in September 2011. The project focuses on transforming the analysis process to provide real-time situational understanding by augmenting the current signature-based analysis with distributed machine-learning behavioral analytics that can be applied as data travels through the network. SETAC requires considerable research in hardware that is powerful enough to apply analytic algorithms to multiple gigabit-per-second data streams and new analytic algorithms that can provide a high-fidelity, real-time behavioral view of the network through distributed agents.

Progress in 2009

SETAC's system capabilities will include (1) identification and classification of high-volume streaming network traffic, (2) high-performance sensor nodes to increase processing capabilities and provide distributed detection across the network, (3) a distributed computing test bed to develop and prototype advanced global collection and analysis approaches, (4) new analysis algorithms to analyze data across the distributed system, and (5) system- and component-level metrics for performance assessment.

As government and military operations become increasingly dependent on computing networks and the Internet, cyber security has grown from a localized economic problem to a major national security issue. SETAC is essential to LLNL's involvement in the Comprehensive National Cyber Security Initiative, the single largest security research and development investment in the country. SETAC is helping explore and establish a vision of 21st-century network awareness sought by the intelligence community and the Department of Defense.

Situational awareness of computer networks presents many challenges because of the scale and complexity of the problem. For example, at the perimeter of LLNL's institutional computer networks, terabytes of network traffic are collected daily, containing millions of unique IP addresses and hundreds of millions of connection records. The cyber threat is extremely dynamic: the nature of the Internet connectivity (topology) changes on the scale of minutes, adversaries employ ephemeral IP addresses to conduct their operations, and polymorphic (constantly changing) malware is prevalent. The complexity of the problem is further compounded by the increasing sophistication of the adversaries' actions, the evolution of the Web, and cutting-edge computing environments that blur the lines between operating systems, applications, personal computers, and networks. All these phenomena contribute to thwarting conventional threat mitigation efforts. Thus, the need is urgent for research that not only delivers heightened awareness for the current computer networks and Internet but also provides viable solutions for the increased scale and complexity of digital networks of the future.

SETAC is a sea change in network situational awareness. It will enable real-time, interactive analysis of the network. Characterizing network-wide activity depends on understanding time-varying patterns of system behaviors, such as the actions and connections between components. Such dynamic relational information will provide a significant advance in the state of the art, but achieving it will require leading-edge data collection, new approaches to building and accessing large-scale graph representations, new capabilities in distributed reasoning infrastructure, and dramatic advances in real-time continuous machine learning that can discover complex time-varying activity patterns.

In 2009, SETAC established a multidisciplinary research team, entered into numerous collaborations and subcontracts, and achieved several milestones. The team developed an agent framework using JADE, an open-source system, and deployed sensor agents to approximately 50 hosts within LLNL. These sensor agents collect host process and network traffic data for analysis. In addition, the team pursued the use of embedded architectures within the network by evaluating two classes of algorithms on the Tilera many-core hardware processing system and on a field-programmable gate array system. SETAC researchers developed and evaluated several behavior-based machine-learning algorithms for cyber security, including HTTP attack detection, network traffic classification, and nonparametric generative models of IP communications. They developed two approaches for finding collections of important hosts (i.e., hosts that cause the most harm if infected) and conducted an

CONTACT

CELESTE MATARAZZO,
MATARAZZO1@LLNL.GOV

extensive comparative study of various bi-flow extractors. The team also developed activity models for LLNL networks, using both network traffic and host process data, reviewed existing cyber security evaluation environments, and began creating a methodology to measure their system's success. The SETAC team co-hosted (with Sandia and Pacific Northwest national laboratories) a grassroots cyber security community meeting to discuss quantifiable scientific evaluation for cyber security, and they worked with the operational cyber security organization at LLNL to gain access to both the local cyber security experts and the data sets required for conducting SETAC research. The team's efforts produced seven peer-reviewed publications with three submissions under review. In addition, they obtained permission to evaluate SETAC on LLNL's institutional network, which will allow them to build performance models that will serve as the basis for scaling to larger sensor and analysis architectures.

In 2010, the SETAC Project will expand the data collection and testing environments by seeking additional volunteers to run the SETAC sensor agent on their desktop machines. This expansion will provide additional opportunities for the sensor agents to look for anomalous scenarios, such as multiple machines simultaneously performing the same action or an unauthorized action.

Cyber security falls within a class of important national and global security challenges that demand improvements in analytical capabilities beyond what are presently available. Current technology is heavily reliant on matching patterns against a human-generated blacklist of known attack features. This approach will never achieve the level of situational awareness needed to characterize and respond to next-generation cyber threats. The nation must prepare for future threats by transforming current paradigms of cyber defense into an approach based on behavioral monitoring, modeling, and analysis of information collected across a broad network. SETAC is addressing these challenges by developing transformational analytic paradigms for efficient, adaptive, and decentralized analysis and computation.

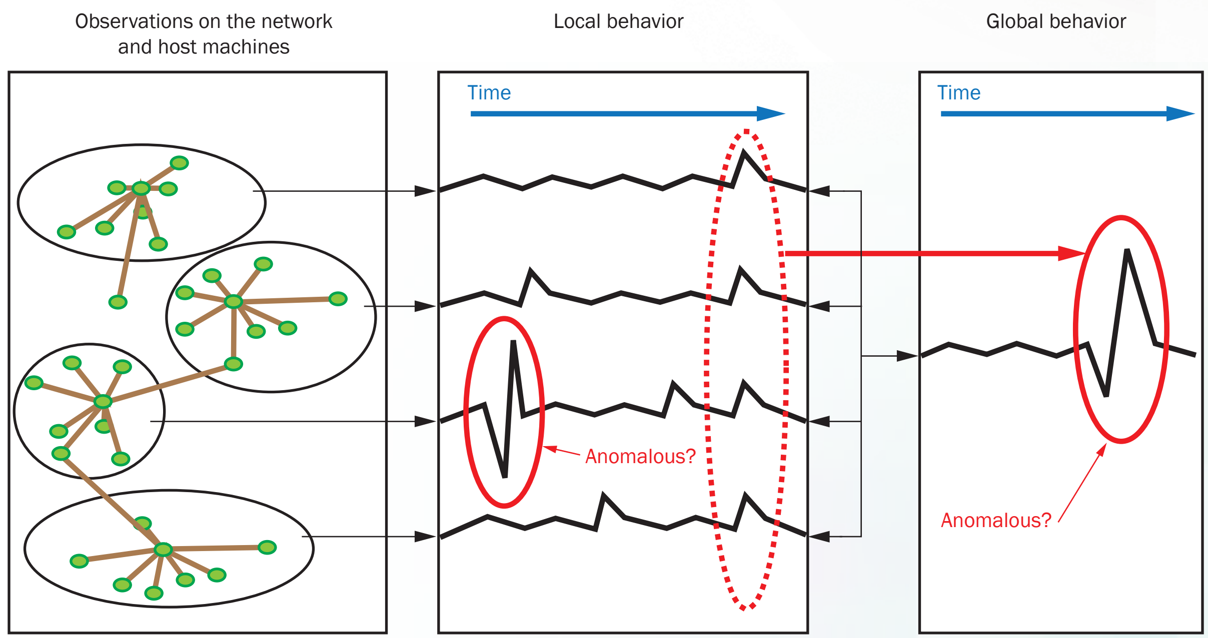

SETAC is designed to detect both local and global patterns of behavior. In this example, sensors continuously observe the behavior of four computers. A spike in activity of one computer indicates a possible anomalous local behavior. A simultaneous spike in activity of all four computers indicates a possible anomalous global behavior. The information is then shared with other sensors, and a security analyst is alerted.

RETURN TO TABLE OF CONTENTS

Space Situational Awareness Helps Safeguard Space Operations

Contact

TAMARA DAHLGREN,
DAHLGREN1@LLNL.GOV

JAMES LEEK,
LEEK2@LLNL.GOV

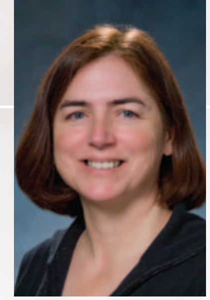

DAHLGREN

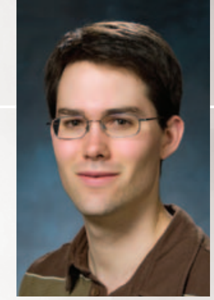

LEEK

Summary

Satellites orbiting the Earth are critical components of a rapidly growing worldwide information infrastructure supporting modern capabilities for a wide range of services, including communications, navigation, and research. Space Situational Awareness (SSA) supports understanding the activities and processes that could affect space operations by providing a fundamental basis for the protection of space assets and helping maintain the delivery of their important services. The rapidly increasing use of space necessitates a corresponding improvement in SSA capabilities. To support this growing need, LLNL is developing a comprehensive SSA modeling and simulation capability and, in collaboration with Los Alamos and Sandia national laboratories, is also working to deploy additional capabilities for collecting and analyzing SSA information in an integrated, net-centric environment using emerging Service Oriented Architecture standards.

Progress in 2009

In 2009, LLNL focused on developing a comprehensive SSA modeling and simulation framework called the Testbed Environment for Space Situational Awareness (TESSA). TESSA is based on a flexible, scalable architecture and exploits LLNL high-performance computing capabilities to enable efficient, physics-based simulation of the current SSA enterprise and to accommodate future SSA systems. TESSA includes hydrodynamic models of satellite intercepts and debris generation, orbital propagation algorithms, radar cross-section calculations, optical brightness calculations, radar system models, optical system models, object detection algorithms, orbit determination algorithms, simulation analysis, and visualization tools. Recent improvements include highly accurate least-squares orbital refinement, massively parallel collision detection, and probability determination. Current efforts are focused on integrating TESSA support for collision debris generation. Theoretical technologies, such as a space-based sensor network, can also be evaluated with TESSA.

LLNL partnered with Los Alamos and Sandia to leverage existing technology in a new Service Oriented Architecture (SOA) to provide services for Air Force personnel to protect satellites and predict the results of orbital events, such as the 2009 Cosmos–Iridium collision. Each laboratory will publish its services over a wide-area network. The SOA allows stakeholders to merge their best-of-breed simulation capabilities and build new simulations using real data. SOA technologies eliminate the necessity of locally recompiling, linking, and maintaining code versions.

An initial set of SOA capabilities has been identified and progress has been made to integrate those capabilities in a coordinated demonstration of tri-laboratory services. Livermore's SOA contributions include debris generation, orbital propagation algorithms, orbit determination from ground-based and space-based sensor measurements, and Monte Carlo collision probability determination. DOE/NNSA high-performance computing, modeling, simulation, and knowledge tools applied in an integrated, net-centric environment will add content and context to the current SSA capabilities using emerging Air Force SOA standards.

This work will result in a fundamentally new capability for SSA that optimizes existing data sources and provides coherent planning for new sensor systems tuned to specific threat scenarios. LLNL's SSA efforts respond to the priorities of the latest national strategies for ensuring the freedom of U.S. space operations and support the core LLNL missions in national and global security.

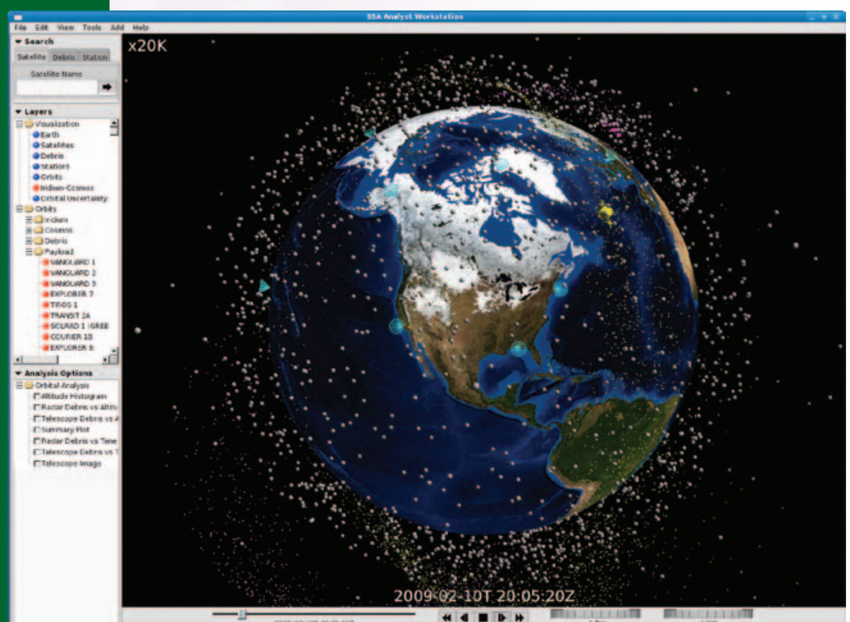

A scalable visualization system for exploring large-scale and time-dependent TESSA simulation data. Gray spheres represent active satellites, yellow spheres Cosmos 2251 debris, magenta spheres Iridium 33 debris, and cyan cones radar stations.

RETURN TO TABLE OF CONTENTS

Protein Structure Analysis Tools Provide Insight into Pathogenic Organisms

CONTACT

CAROL ZHOU,
ZHOU4@LLNL.GOV

ADAM ZEMLA,
ZEMLA1@LLNL.GOV

AS2TS.LLNL.GOV

ZHOU

ZEMLA

SUMMARY

Computation scientists are developing a suite of modeling and analysis tools for elucidating the structures and functions of proteins. Their work is motivated by the need for improved bioinformatics methods for understanding bacterial and viral genomes, including those of emerging, engineered, and advanced threat pathogens. A whole-genome, computational approach to pathogen characterization requires multiscale modeling that bridges the various scales of biological organization, from systems biology to mechanistic modeling. The accurate protein-structure-to-function predictions being developed in the Pathogen Bioinformatics group at LLNL are a middle ground within this multiscale approach. In anticipation that structure-to-function methods would eventually be applied in high-throughput for applications in pathogen detection and medical countermeasures, researchers have directed their efforts toward devising novel algorithms for identification of structure–function motifs and quantification of amino-acid residue frequency in structure context, which are of particular importance in understanding rapidly evolving ribonucleic acid (RNA) viruses.

PROGRESS IN 2009

LLNL has been recognized as a leader in protein structure analysis since 1996, when scientists designed and established the first fully automated protein structure prediction evaluation system (ACE) to support the international Critical Assessment of Techniques for Protein Structure Prediction (CASP) competition. Novel numerical criteria for detecting protein structure similarities, as implemented in the Local-Global Alignment tool, became and remain the standard structure evaluation approach for assessing the accuracy of structural models submitted to CASP. For the past 14 years, LLNL scientists have continued to design and develop structure-based tools for protein structure-to-function prediction with the goal of serving LLNL's biodefense mission in pathogen identification and characterization and ultimately contributing to the development of advanced medical countermeasures.

In 2009, the Computation team designed and implemented a novel algorithm for quantifying protein sequence variability in structure context and applied the algorithm in a study of the RNA-dependent RNA polymerase protein from poliovirus, the causative agent of poliomyelitis. Poliovirus is considered a quasispecies because it exists as a "cloud" of diverse but closely related genomes centered around a dominant type. The rapid evolution of quasispecies viruses makes managing chronic infections (e.g., AIDS) and controlling emerging infectious diseases (e.g., SARS) very difficult. Because poliovirus has been well studied, it serves as a good model system for the validation of new bioinformatics tools.

The team's newest code, Structure-Alignment-Based Sequence Variability (StralSV), uses structure-based alignments and leverages data from all experimentally determined protein structures to quantify the amino acid residue variability at each position in a protein of interest (the reference protein). StralSV searches the Protein Databank database of experimentally determined protein structures and selects structurally similar fragments, aligns them along the reference protein, and quantifies how many different amino acids occur at corresponding positions in the fragments. Data are collected to determine how many fragments contributed to the variability statistic at each position. Subsequent analyses enable automated identification of structure and sequence motifs that are common among proteins with similar structure and function and help generate hypotheses regarding protein function.

This work supports an ongoing effort to construct a protein-centric bioinformatics capability that will meet the needs of the biology, medical, biodefense, and biofuels research communities, and attract support from agencies including DOE, the Defense Threat Reduction Agency, and the Department of Homeland Security.

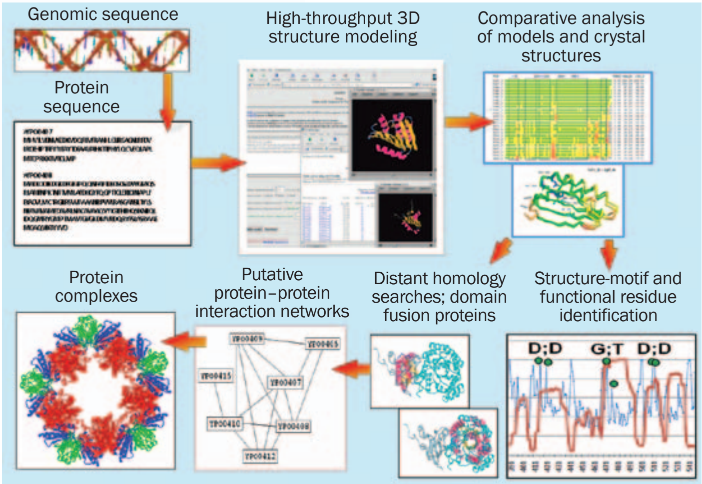

Overview of LLNL structure-based bioinformatics tools for characterizing protein structure and function.

RETURN TO TABLE OF CONTENTS

Parallel Streams Catapult Data Analysis to the Petascale

Summary

With scientific data sets from observation and simulation reaching peta- and even exabytes in the near future, our data analysis capabilities are being outstripped by our ability to generate such enormous volumes of data. This trend is compounded by the increasing mismatch between plentiful compute power available on today's multicore processors and low disk and memory access speed, which are improving at a much slower pace. This performance gap is causing data transfer to be the bottleneck both in simulation codes and in the subsequent analysis of the data from distributed storage.

Computation scientists are developing new techniques to mitigate the data transfer bottleneck in two complementary ways. First, they use custom compression algorithms specialized for scientific data to reduce the amount of data that needs to be transferred and stored, and thus trade idle processor cycles for more effective bandwidth. Second, they minimize the number of slow data accesses via single-pass streaming computations and reorganize data and computations to improve locality of reference and better use fast caches. Computation researchers have used these techniques to demonstrate data analysis of unprecedented scale and speed.

Progress in 2009

Data analysis and visualization are fundamental to gain insight into the complex physical phenomena being simulated at extreme resolution and scale at LLNL (e.g., climate prediction, inertial confinement fusion, and weapons performance). Current approaches to large-scale data analysis are nearly without exception limited to parallel in-memory solutions that rely on spatial domain decomposition of the data set into pieces that are small enough to fit in memory on a compute node. Collective communication is then used to share partial analysis results and resolve data dependencies across domain boundaries. Until recently, this model of brute-force parallelism has been successfully employed by visualization and analysis tools like LLNL's VisIt, which relies on a dedicated analysis cluster with large enough aggregate main memory to hold and process data sets from some of the world's largest simulations. However, on next-generation supercomputers, such as the ASC Sequoia machine with an expected 1.6 petabytes of random access memory, the amount of data generated is expected to dwarf the memory of even the most powerful analysis clusters, rendering this in-memory approach untenable in the long run.

Over the past several years, Computation's research team has developed a framework for out-of-core data processing based on the concept of windowed streaming. This approach treats data, such as a two- or three-dimensional field defined on an unstructured computational mesh, as a one-dimensional linear stream of data elements that pass sequentially through main memory on a single compute node with only a small portion of the data set visible to the computation at any one point. A series of streaming tasks can be chained together in a dataflow pipeline, much akin to an assembly line, to achieve task parallelism. Moreover, since streaming input/output (I/O) is sequential, data compression is easily integrated and hidden from the analysis application.

Reading and processing a petabyte of data on a single compute node using today's disk technology would take years to complete. With out-of-core computing still a necessity and parallel analysis clusters readily available, Computation's recent efforts have targeted stream computing on parallel computers. The team's approach is to spatially organize and partition the data into small "streamlets"—blocks of data that can be streamed and processed atomically—and to surround each block with one or more layers of "ghost elements" from adjacent blocks. This duplication of data between overlapping ghost layers allows decoupling data dependencies between adjacent blocks such that the analysis of each block can proceed independently. For instance, a gradient computation requires information from the immediate neighbors

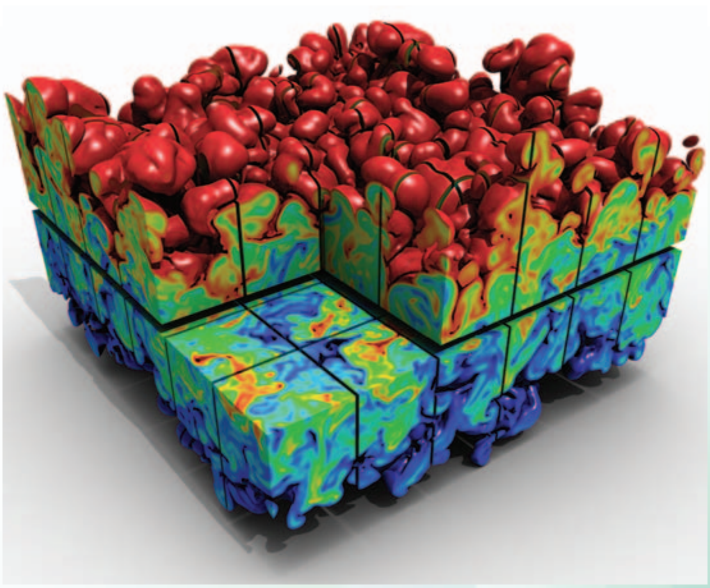

A structured volumetric grid is partitioned into independent blocks that can be streamed and analyzed in parallel.

Contact

PETER LINDSTROM,
PL@LLNL.GOV

COMPUTATION.LLNL.GOV/CASC/LOCAL/

of each grid element and can thus be implemented by first attaching to the block one layer of elements from adjacent blocks as ghost elements. The novel aspect of this ghost layering approach is that it can be accomplished not only in parallel but also as an out-of-core streaming computation. One at a time, data blocks assigned to a processor are padded with ghost elements before being passed on to a sequence of one or more downstream analysis tasks. The result is a flexible and efficient framework that allows both task and data parallelism at any desired granularity.

In contrast with the nonstreaming, in-memory ghost layer computation previously employed in VisIt, only one or a few blocks per processor reside in memory at any point in time. This method allows the computation to finish on any number of processors—possibly just one—and use whichever compute resources are available. Moreover, this preprocessing step fits directly into VisIt's dataflow paradigm, allowing it to be integrated with only minor code changes. The new technology has enabled VisIt to analyze data sets of unprecedented size, including grids with trillions of elements, using only modest amounts of main memory. The streaming parallel approach requires only a small, constant amount of memory per processor regardless of the number of processors involved in the computation, whereas VisIt's previous nonstreaming approach required a minimum of 128 processors and nearly all the aggregate memory to complete. Another benefit of a fully streaming approach is that, in spite of running out of core, it executes 30% faster than VisIt's in-memory method, in part due to overlapped computation and I/O, less interprocessor communication, and better cache use. With analysis speed scaling linearly from a single to hundreds of processors, Computation's parallel streaming techniques are an important step toward dealing with the vast amounts of data expected at the petascale and beyond.

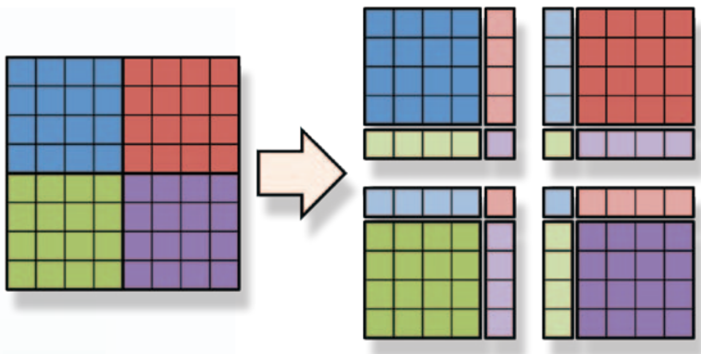

Data dependencies can be eliminated by padding blocks with one row and one column of ghost elements from adjacent blocks.

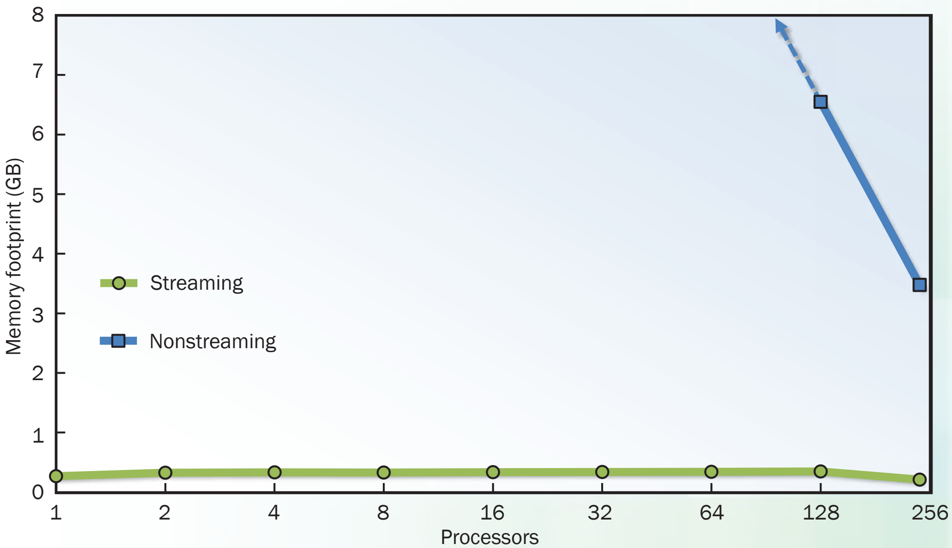

Maximum per-processor memory footprint of an analysis task involving 27 billion elements. This plot illustrates the advantage of the new parallel streaming approach compared to VisIt's previous nonstreaming method.

Innovative Algorithms Extend Laser–Plasma Interaction Simulation Capabilities

CONTACT

JEFFREY HITTINGER,
HITTINGER1@LLNL.GOV

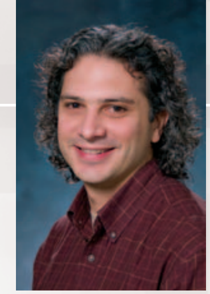

SUMMARY

Future experiments at the National Ignition Facility are expected to enter regimes of high-energy–density physics where nonlinear laser–plasma interactions (LPI) are important. Direct simulation of the Vlasov equation in a high-dimensional particle phase space is a potentially useful approach for certain LPI problems, such as identifying nonlinear thresholds of parametric instabilities. However, the computational costs of this technique have been prohibitively expensive. Computation researchers are leading a collaboration with physicists in the Physical and Life Sciences Directorate and the Weapons and Complex Integration Principal Directorate to develop new, efficient, high-fidelity algorithms that will facilitate routine, direct Vlasov simulations for LPI on massively parallel computers.

PROGRESS IN 2009

The principle difficulty with direct Vlasov equation discretization is that the number of discrete unknowns increases geometrically with dimension. For problems in three spatial dimensions, solutions of the Vlasov equation are defined over a six-dimensional phase space. Typical discretization methods for the Vlasov equation use uniform grids and fail to conserve particle number, control numerical oscillations, or preserve solution positivity. Furthermore, these schemes are not optimal for data and task parallelization. Computation's approach focuses on techniques that reduce the number of discrete unknowns within a framework that promotes conservation, positivity preservation, and parallelization on modern architectures.

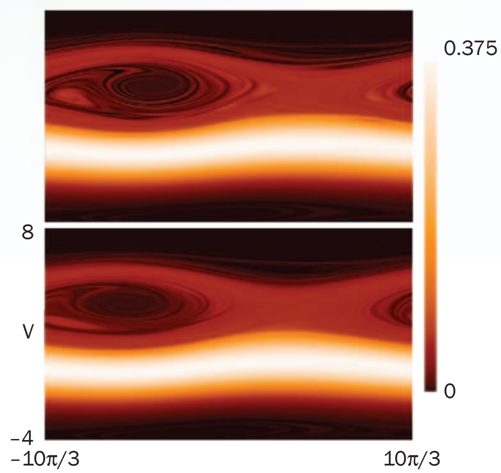

Particle distribution function for a two-dimensional bump-on-tail problem using the new scheme with one million unknowns (top) obtains better-resolved features than a standard limiting scheme using 16 million unknowns (bottom).

Computation researchers developed a new class of weakly limited, high-order finite-volume algorithms that meet these design criteria. High-order methods require fewer unknowns to represent solutions at a fixed level of error, so researchers developed nonlinear flux functions based on ideas from the Weighted Essentially Non-Oscillatory family of schemes. These numerical flux functions for the Vlasov equation are central-differenced, dissipation-free discretizations in well-resolved regions, but they reduce to one-order-lower, upwind-differenced discretizations in under-resolved regions to provide sufficient dissipation to control unphysical oscillations. Researchers implemented a nominally fourth-order flux function with a four-stage, fourth-order Runge–Kutta time integrator in two dimensions and four dimensions. Preliminary comparisons with semi-Lagrangian methods on the Vlasov–Poisson system indicate better performance of the new scheme using fewer unknowns for a comparable level of error. Near-perfect scalability has been achieved on 1,000 processors. Further gains can be realized from task parallelism because the explicit method-of-lines formulation allows for concurrent evaluation of spatial discretization terms.

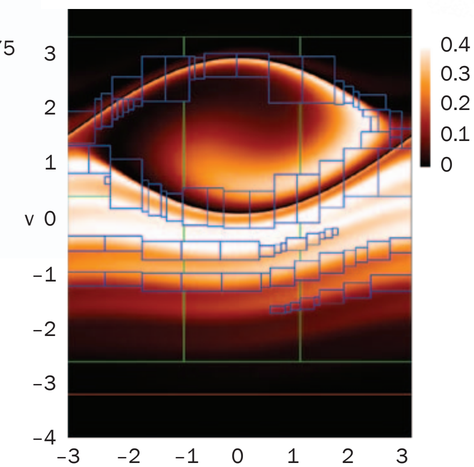

Particle distribution function for a two-dimensional imposed electric field problem. Four levels of AMR reduced the number of unknowns by more than 70%.

The largest computational gains in the researchers' work come from the use of adaptive mesh refinement (AMR), which refines grids only in regions of poorly resolved solution in phase space. They leveraged the block-structured AMR infrastructure of LLNL's SAMRAI library in a new parallel simulation code, VALHALLA. Because the new discretization is in a discretely conservative, flux-divergence form, the algorithmic adaptation for AMR is trivial. The researchers conducted proof-of-principle calculations on a forced Vlasov problem that demonstrated a mesh reduction of more than 70% in two dimensions, and they expect even greater savings in higher dimensions. Computation researchers are completing the infrastructure required to handle parallel, interdimensional data transfers between AMR mesh hierarchies needed to enable the first-ever, fully coupled Vlasov–Poisson, block-structured AMR simulations. Future work will include extending AMR implementations to higher dimensions, generalizing to an electromagnetic model, and enabling multiple kinetic species.

RETURN TO TABLE OF CONTENTS

Applications Software Development

John Grosh,
grosh1@llnl.gov

The Computation Directorate partners with all of the Laboratory's programs and numerous external organizations to develop software used in a variety of scientific and mission-specific applications. Examples include codes for large-scale multiphysics simulations, tools for understanding large data sets, and frameworks for controlling sensors and experiments. To develop these products, Computation has organized more than 400 computer scientists, software engineers, mathematicians, engineers, physicists, and technicians into the Computing and Applications Research (CAR) Department.

CAR Department software-related activities span from basic research in computer science and applied mathematics to development and systems integration using commercial products. Projects are conducted in collaboration with staff in other directorates, academia, industry, and other government laboratories. Software development activities are managed using a strategy of anticipate, innovate, and deliver. Anticipate includes high-risk research used to build new concepts and programs, innovate produces award-winning technologies that provide unique capabilities, and deliver builds usable software focused on end-user requirements.

Multidisciplinary teaming, excellence in software engineering, and algorithm and computer science research expertise are key to developing successful software at Lawrence Livermore National Laboratory (LLNL). The Laboratory has a tradition of building robust usable applications, achieved through teams of domain experts and software developers.

This section includes examples that illustrate these attributes of success. The first is an article on enhancements to ARES, a large multiphysics simulation code critical to the National Nuclear Security Administration's (NNSA) Advanced Simulation and Computing (ASC) Program. The ARES code team addressed two important challenges: improving execution performance, and re-engineering the architecture to improve the extensibility for incorporating new physics and other features. New algorithms and scalable data structures were developed to integrate Arbitrary Lagrangian–Eulerian (ALE) meshing schemes into the SAMRAI adaptive mesh refinement (AMR) library. The integration of ALE and AMR techniques was novel and yielded an order of magnitude improvement in performance. The enhanced ARES code has demonstrated strong scaling to more than 140,000 processors. In addition, software architecture experts re-engineered the ARES architecture, creating a new object-oriented framework for the code. Such successes cannot be achieved by "throwing software over the fence." In the case of ARES, success depended on a close partnership between researchers, software developers, and the physicists who aggressively confronted the integrated problem of performance and extensibility.

The second example of success is the development of system models used to support the planning and operations of the nuclear security weapons complex. This software predicts and conducts trade-off studies to examine the influence of funding, capital investments, policy, and workforce issues on the effective management of the nuclear weapons complex. Compared to the previous examples of software developed mostly by LLNL staff, the Nuclear Enterprise Modeling software required the integration of complex third-party software to provide a solution. The important point is that software solutions can take many forms—from homegrown to commercial—and require a wide range of expertise to ensure their success.

Another example is software developed in support of the National Ignition Facility (NIF) diagnostic systems. The Precision Diagnostics System (PDS) is used to coordinate and control various sensors, data acquisition, and analysis tools to ensure optimal laser beam performance. As NIF continues its National Ignition Campaign, flexibility and adaptability of PDS is required to support various proposed experiments. To this end, Computation researchers developed the Instrument-Based Controller (I-BC) architecture, which is used to program and control the hardware and software that operate diagnostic instruments. The I-BC architecture was developed using Java and CORBA, and the process involved integrating Ada code and other software technologies, such as Hierarchical Data Format. A dozen diagnostic systems were implemented this year using the innovative I-BC architecture.

Future software applications will be used to model, simulate, and analyze systems of increasing complexity, and more attention will be given to the quality of the answers provided by the applications software. As a result, there will be a greater need for tools and uncertainty quantification, verification, and validation. These trends will drive improvements and enhancements that will result in larger and more complex software systems. Thus, expertise in software architecture, a wide range of software technologies, and the ability to form multidisciplinary teams of software developers, researchers, and domain experts will become increasingly important to enable the Laboratory to accomplish its various missions.

RETURN TO TABLE OF CONTENTS

ARES Code Enhancements Improve Parallel Scaling Performance

Summary

ARES is a massively parallel, multiphysics computer code developed at LLNL. It is a principal NNSA/ASC tool used to simulate and analyze a variety of science and engineering problems critical to national security and other Laboratory missions. Such problems include laser-driven, magnetically driven, and high-explosive-driven multimaterial systems.

Many important multiphysics systems that are of interest to LLNL programs involve complex, dynamic, three-dimensional phenomena, such as chemical reaction fronts, shock waves, and multimaterial interactions. Accurate simulation of this behavior often requires an extremely high level of model fidelity and numerical resolution that pushes the limits of modern numerical algorithms and available computational resources. In 2009, Computation scientists advanced the capacity of the ARES code to simulate larger and more complex problems. Two key emphases of this effort involved introducing AMR functionality and improving parallel scaling performance.

Progress in 2009

The ARES developments are rooted in two research and development activities in Computation's Center for Applied Scientific Computing (CASC). The first effort is the Structured Adaptive Mesh Refinement Application Infrastructure (SAMRAI) software library, which has been under development in CASC for more than a decade. SAMRAI has served as an AMR algorithm research and application development platform for a wide range of projects both internal and external to LLNL. In the second effort, CASC researchers used SAMRAI to develop a hybrid hydrodynamics capability that combines AMR with ALE methods, which are employed in ARES. The recent parallel scaling performance improvements achieved in ARES were the result of replacing the parallel data communication software in the code with a new, more powerful infrastructure that is based on the SAMRAI software and algorithm design.

ALE hydrodynamics methods, like those used in ARES, combine desirable features of Lagrangian methods, suited to moving boundaries and multiple materials, and Eulerian methods, suited to strong shearing and vortical motion. The result is a method that applies to various problems involving multiple materials interacting in complex ways through their respective motions. However, the computational resources required to achieve adequately resolved details for some problems make their simulation prohibitive and sometimes intractable. AMR is a numerical technique that provides automatic mechanisms for dynamically changing numerical mesh resolution locally. This capability allows computational resources to be concentrated near features (e.g., material boundaries and reaction fronts) that require enhanced detail without incurring the excessive cost of over-resolving other problem regions away from these features.

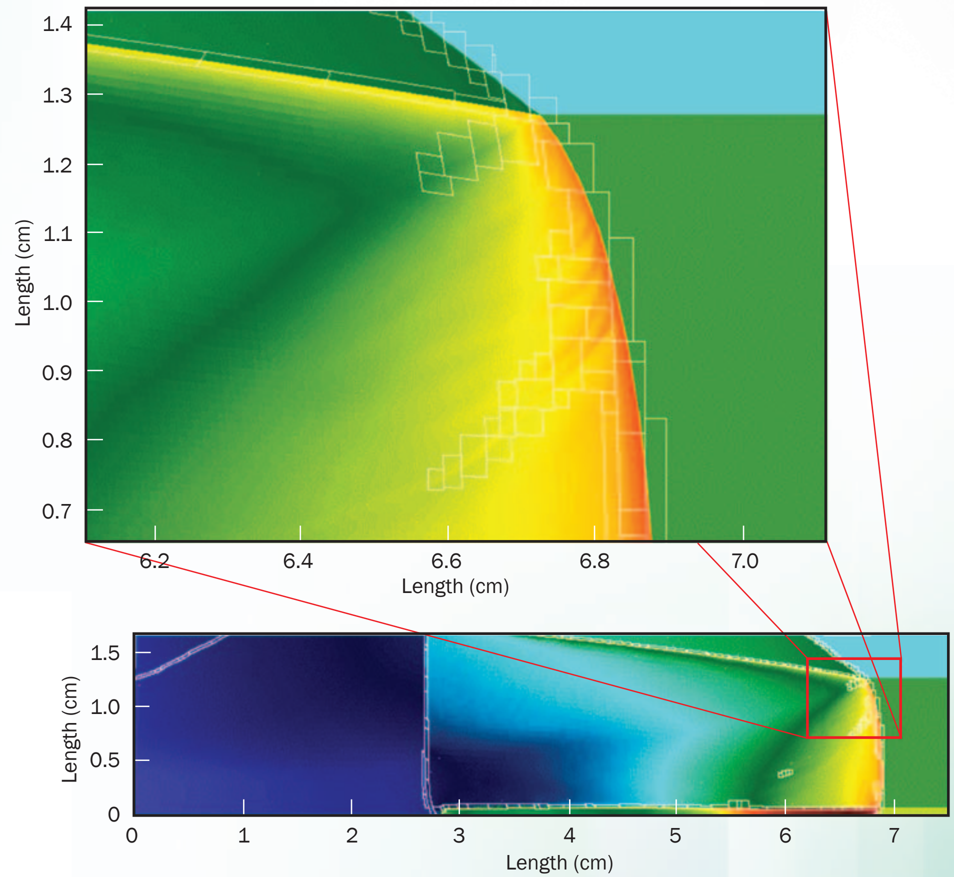

The significant details of fine-scale reacting flow chemistry are captured very efficiently using ALE–AMR algorithms in ARES. The white outlines show the finest mesh regions where fine-scale physics resolution is required for accurate simulation.

Contact

Richard Hornung,
hornung1@llnl.gov

The hybrid ALE–AMR methodology developed by CASC researchers preserves the moving, deforming meshes of ALE methods that are needed to capture material interfaces and provides the flexibility to dynamically add and remove mesh points. ALE–AMR has proved to be a powerful technique for increasing the efficiency and accuracy of multimaterial hydrodynamic simulations. It has been shown to yield order-of-magnitude performance gains in achieving simulations with a given accuracy. More importantly, it has enabled the exploration of ultrahigh-resolution regimes that would otherwise be intractable on some computing platforms. The figure on page 28 shows how ALE–AMR in ARES allows grid refinement to be concentrated near fine-scale features in a reacting flow calculation. In 2009, essential elements of ALE–AMR algorithms and SAMRAI software were integrated into the main line of ARES code development. Computation researchers are currently restructuring the code to streamline this integration, facilitate the adoption of AMR functionality by ARES users, and expand the range of AMR-aware physics in ARES. A production-level AMR capability will be available to users by late 2010.

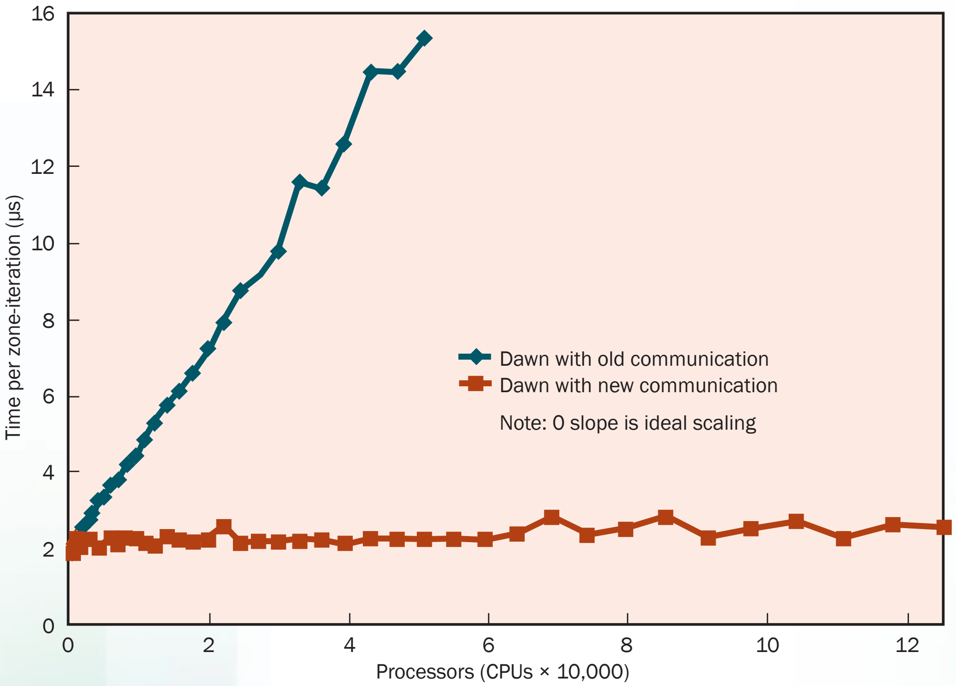

The new ARES communication infrastructure (red curve) shows vastly improved weak scaling over the old routines (blue curve) for a three-dimensional radiation diffusion problem on Dawn. Also, the reduced memory footprint of the code at large-scale enables it to run on much larger processor counts than before. The blue curve ends where the old version of the code ran out of memory.

In 2009, researchers also significantly improved the parallel scaling performance of ARES by replacing the parallel data communication software in the code with a new, more powerful infrastructure. The figure at left shows a weak scaling result on a radiation diffusion problem run at full scale (147,456 cores) on the ASC Dawn system. The new infrastructure scales much better than its predecessor. The scaling shown in the figure is almost identical to ARES' scaling on the ASC Purple machine but at processor counts far beyond what Purple has. Also, the code's memory usage scales significantly better with the new infrastructure. By using far less memory at large scale than the original code, ARES can run on many more processors than it could before.

Beyond performance advantages, the new object-oriented communication software simplifies the addition of new physics capabilities in the code by making it easier for developers to combine existing software entities in novel ways. Also, since its design is based on that of SAMRAI, the new code better integrates ALE–AMR capabilities.

Parallel performance work in the ARES code currently involves exploring shared memory parallelism and AMR performance. While the SAMRAI infrastructure has demonstrated weak scaling of AMR operations on the Dawn machine beyond 36,000 processors, ARES–AMR operations require further analysis to run problems efficiently on more processors.

Also in 2009, shared memory parallelism capabilities were reworked so that ARES can flexibly run multiple threads and/or multiple Message Passing Interface (MPI) tasks per node, depending on the architecture. This work was driven by the desire to take full advantage of shared memory characteristics in new multicore architectures. Preliminary studies have shown the feasibility of the implementation, but a significant amount of work remains to reap the potential performance advantages of a mixed MPI/thread programming model.

Nuclear Enterprise Modeling Helps Envision the Future

CONTACT

MARILYN PICKENS,
PICKENS1@LLNL.GOV

JEENE VILLANUEVA,
VILLANUEVA6@LLNL.GOV

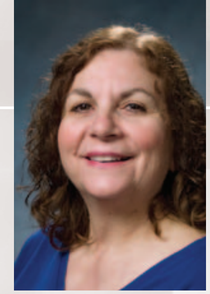

PICKENS

VILLANUEVA

SUMMARY

LLNL's enterprise modeling efforts help Department of Energy (DOE) and Department of Defense (DoD) leaders reach informed decisions about the future state of the U.S. nuclear security enterprise. Recent administrations are pushing toward zero deployed nuclear weapons, but the need to maintain the nation's nuclear deterrent persists. Policymakers must be responsive to this push while taking into consideration the complex realities of a vast nuclear enterprise. Some of these considerations include: how to maintain deterrent capabilities considering the age of weapons and delivery systems, the capacity to refurbish equipment, and uncertain funding models; how to sustain or consolidate facilities while accounting for associated costs and risks and the effects on infrastructure and personnel; and how to map the future composition of weapons scientists and engineers, taking into account demographics, critical skill sets, retention, recruitment, and budgets.

To properly evaluate these scenarios, policymakers and leaders need a computational modeling system that integrates diverse sources of data, builds analytical tools to project various multidimensional courses of action, and provides graphical interpretation. Enterprise modeling requires a diverse team with expertise in databases, simulation, numerical methods, graph theory, and user interface design.

PROGRESS IN 2009

Computation's Nuclear Enterprise Modeling (NEM) team supports stockpile lifecycle, infrastructure transformation, workforce simulations, and policy and budget impact projections for several customers. Computation's enterprise modeling suite of applications includes custom tools that interface with commercial off-the-shelf (COTS) software for advanced simulation and analysis in areas such as systems dynamics, optimization procedures, probabilistic risk assessment, Monte Carlo statistical methods, discrete event simulation, automated graph drawing and network analysis, and multi-attribute decision tools.

Computation's tools are unified by a nuclear enterprise modeling database and application, called NEMUS, which consolidates and organizes the disparate NEM data streams (stockpile, budget, U.S. policy and program objectives, infrastructure, program execution, and manpower/critical skills) using the industry standard tools, Oracle and Java. Combinatorial graph theory is used for workforce modeling, which involves understanding how employee attrition and institutional hiring policies impact the intellectual resources needed to execute NNSA's Stockpile Stewardship Program.

In 2009, Computation collaborated with DoD on Strategema, a custom software product that integrates defense policymaker options with concrete models of U.S. nuclear force structure, stockpile dynamics, and nuclear enterprise infrastructure in response to policy, military, scientific, and budgetary drivers. Also in 2009, LLNL was designated as the Chair of the Enterprise Modeling Consortium—an extensive collaboration with the eight NNSA Defense Program sites. In addition to new development efforts with the Defense Threat Reduction Agency, Johns Hopkins University, and the Air Force Nuclear Weapons Center, Computation scientists are collaborating with the United States Strategic Command, the Office of the Secretary of Defense, and industrial and academic partners. LLNL's modeling tools will be critical for evaluating options in the post-Nuclear Posture Review nuclear security enterprise.

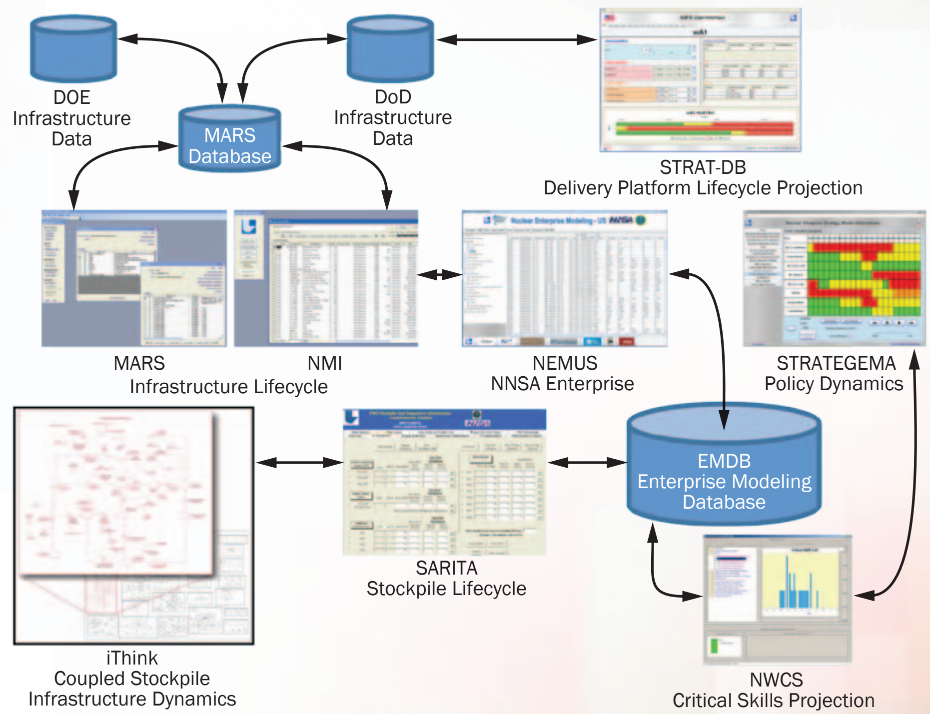

Enterprise modeling incorporates various COTS software and custom codes for calculation and analysis. The NEM team uses an array of tools to provide insight to the Nuclear Posture Review working groups.

RETURN TO TABLE OF CONTENTS

Anticipating Mesh Partitioning Needs on Petascale Machines

CONTACT

KATIE LEWIS,
LEWIS66@LLNL.GOV

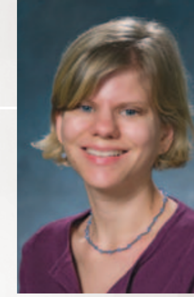

SUMMARY

The Parallel Mesh Generation (PMESH) Project is addressing critical issues in exascale computing by identifying deficiencies that are already problematic for petascale machines. Mesh partitioners divide a mesh into submeshes called domains; simulations run in parallel by assigning each domain to a different processor. Partitioners attempt to divide work evenly among the processors (load balancing) while minimizing interprocessor communication (represented by edge cuts in graph algorithms). Experiments show that existing parallel mesh partitioning algorithms will not produce adequate results on complex meshes at the petascale and that many partitioners fail to make adequate numbers of domains. Finding and addressing inadequacies with current mesh partitioners is crucial to running accurate three-dimensional simulations, a goal essential to DOE and the high-performance computing community.

PROGRESS IN 2009

Computation's PMESH team ran strong and weak scaling tests on four parallel partitioners: (1) ParMETIS from the University of Minnesota; (2) a recursive coordinate bisection (RCB) algorithm from LLNL; (3) a PMESH algorithm that creates block-structured domains; and (4) Parallel Hypergraph (PHG) from Sandia National Laboratories. As part of the tests, the team varied the algorithms' quality control parameters to the extent possible. They used a realistic mesh that stresses algorithms in several ways: some mesh zones are very small and have high aspect ratios, some have many blocks with widely varying numbers of zones, the geometry has nonconvex regions, and the mesh has many points of enhanced and reduced connectivity.

The team's strong scaling study divided an approximately three-million zone mesh into 2n domains for every power of 2, from 512 to 131,072. The weak scaling study created the same set of domains while adjusting the mesh size to maintain a constant average domain size of approximately 3,000 zones. The team's partition quality metrics extended beyond the standard edge-cut metric to reflect the specific needs of the simulation codes Computation supports, using information about the codes' communication schemes to more accurately estimate the communication costs.

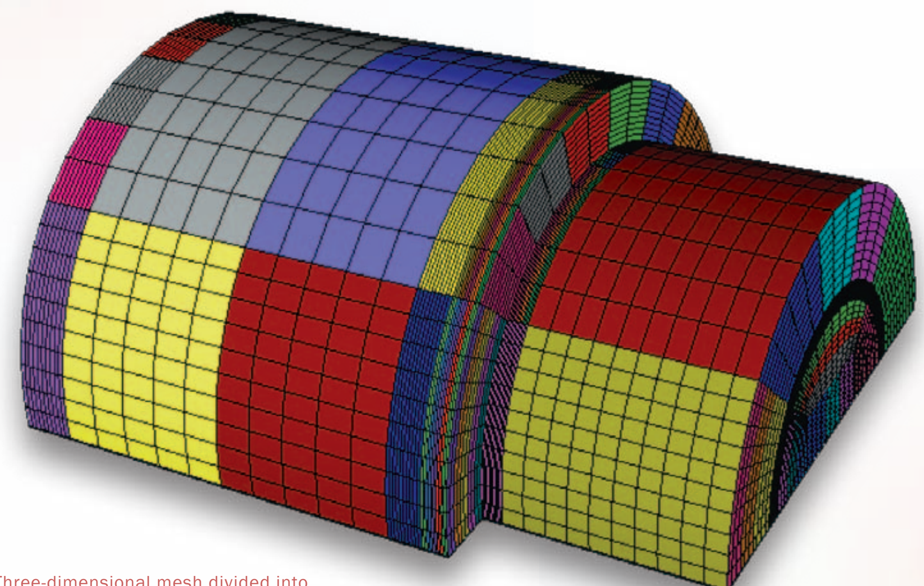

Three-dimensional mesh divided into 512 block-structured domains.

The results indicate that a good solution for mesh partitioning remains elusive on petascale systems. ParMETIS failed to make more than 32,000 domains, even in the strong scaling tests. RCB and the block-structured algorithm reliably created large numbers of domains, but partition quality was often poor. RCB was also prone to poor communication because it uses the geometry but not the topology of the mesh. The block-structured approach produced poor load balancing because it must create topologically rectangular domains lying within unequally sized mesh blocks. On large numbers of domains, PHG created domains with zero zones, which may cause simulation codes to fail. Empty domains existed even for the weak scaling cases, where domains averaged 3,000 zones. The PMESH team does not have the full set of results for PHG, due to its memory usage on large meshes.

The PMESH Project will continue testing how mesh partitioners scale by evaluating the PT-Scotch partitioner and a partitioner being developed in CASC by the Sparta Project. The team is working with the mesh partitioning community, including scientists at Commissariat à l'Énergie Atomique, to achieve the partitioning capability necessary to fulfill critical DOE missions.

RETURN TO TABLE OF CONTENTS

Computation Supports NIF's March toward Ignition

Summary

NIF was completed in 2009, and the world is watching LLNL's progress toward achieving fusion ignition. This year's initial National Ignition Campaign (NIC) experiments demonstrated that the program is poised to achieve ignition within the next two years. If successful, NIF will open the door to profound advancements in energy, defense, and astrophysics.

NIF's accomplishments are a result of more than a thousand employees from almost every organization at LLNL and hundreds of partners and vendors around the world who helped build and are now conducting experiments on the world's largest laser facility. More than 100 Computation employees are part of NIF's multidisciplinary team. They are primarily focused on delivering NIF's state-of-the-art controls, diagnostics, shot analysis, and modeling and simulation aspects. This year marked the delivery of a mature set of target diagnostics systems and the computational foundation for a fully integrated NIF experimental planning environment.

Progress in 2009

Twelve NIF target diagnostics subsystems were deployed in 2009 to provide data on energetics, laser–hohlraum interaction, hydrodynamics, and equations of state of materials for NIF energetics experiments. Each diagnostic has unique controls requirements; some collect data at nanosecond and picosecond timescales. Optical diagnostics observe backscattered light from targets and provide insight into energy conversion. X-ray diagnostics can be either integrating, time resolved, or gated to capture a snapshot of the target experiencing laser-driven shock. Neutron time-of-flight and spectroscopy are used in diagnosing ignition experiments.

NIF's Integrated Computer Control System (ICCS) was enhanced with controllers for the 12 new diagnostics using the Instrument-Based Controller (I-BC) architecture. The instrument-based architecture couples inexpensive controller hardware with modular configurable software to control a single instrument within a diagnostic. For example, a diagnostic's supporting instruments (i.e., power supplies, cameras, and/or digitizers) are each supported by a dedicated I-BC with software specific to that instrument. This innovative architecture has reduced development and deployment time for new or modified diagnostics.

The controller software was developed using the I-BC framework, an objected-oriented, Java-based, modular, data-driven software library that provides all the major functions necessary to create controls software for specific target diagnostics instruments. The I-BC framework is composed of objects that perform functions such as reading Extensible Markup Language configuration and setup files, creating Hierarchical Data Format data archive files, migrating data set management, automatically generating remote graphical user interfaces, and providing access to device drivers. In addition, the objects implement the Diagnostic Control Protocol (DCP) that loosely couples each I-BC device to ICCS.

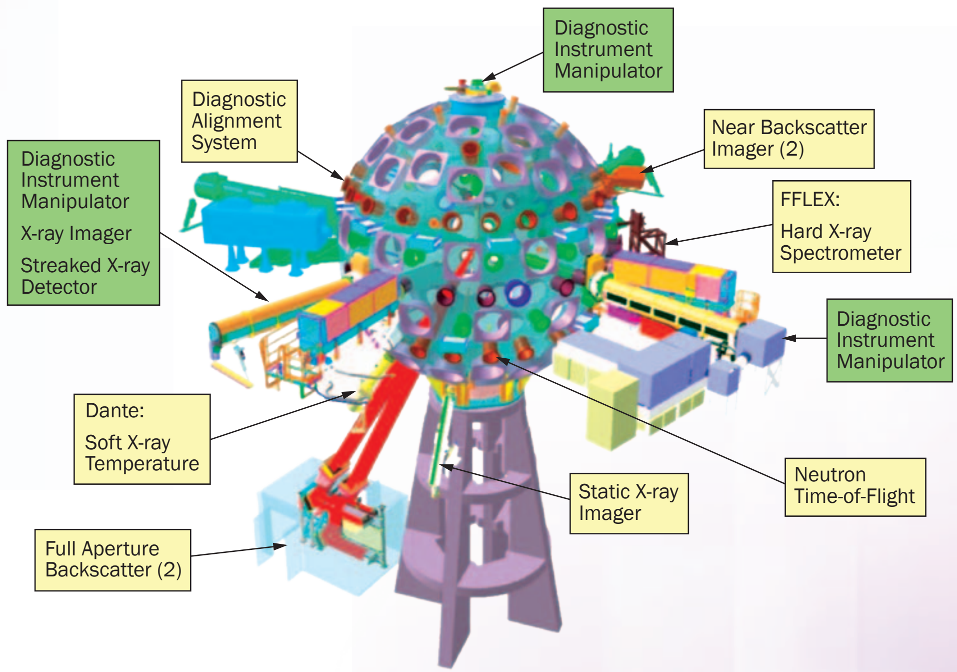

NIF target diagnostics with 200 data channels have been activated for the energetics experiments.

Contact

RANDY SHELTON,
SHELTON1@LLNL.GOV

JOSEPH TAPPERO,
TAPPERO1@LLNL.GOV

SHELTON

TAPPERO

The target diagnostic subsystem software is written in Ada and runs on Sun Solaris servers. It is composed of three modules. The Diagnostic Bridges module translates the DCP from each I-BC into ICCS Common Object Request Broker Architecture objects. The Target Diagnostics Supervisor module uses these bridges to provide status and control of each controller and groups them for the diagnostics they support. It also provides the primary operator interface on the target diagnostic console in the NIF control room. The Target Diagnostics Shot Supervisor module runs automated sequences that are defined by the engineer responsible for a particular diagnostic. Because of this extensible architecture, NIC fielded a dozen diagnostics systems and is poised to rapidly deploy a dozen more diagnostics in 2010.

In October 2009, the foundational components of the NIF Shot Planning and Analysis Tools (SPLAT) Project were deployed. The SPLAT Project is the result of forward-thinking operational concepts that former NIF Chief Technology Officer Mary Spaeth developed several years ago. The project combines near- and long-term shot planning, final optics demand/supply loops, target diagnostics planning, and target fabrication requirements. The October deployment was the culmination of months of fulfilling requirements, creating systems architecture, designing applications, and implementing and testing the foundational components. The primary focus was on the shot planning and optics demand components; the next step is to gather requirements for the target diagnostics and target fabrication planning components. The long-term goal of the SPLAT Project is to integrate the components that will improve NIF's overall operational efficiency.

SPLAT's shot-planning component provides a state-of-the-art Web-based interface for selecting and building a sequence of proposed shots for NIF. These shot sequences, or "lanes," help plan the near-term shots in the facility and the long-term campaigns. The architectural shot-planning foundation deployed this year integrates with the established tools for experimental configuration management that specify details of the experiment, as well as the NIF calendar of availability, and it will eventually integrate with the target diagnostics planning and target fabrication requirements tools.

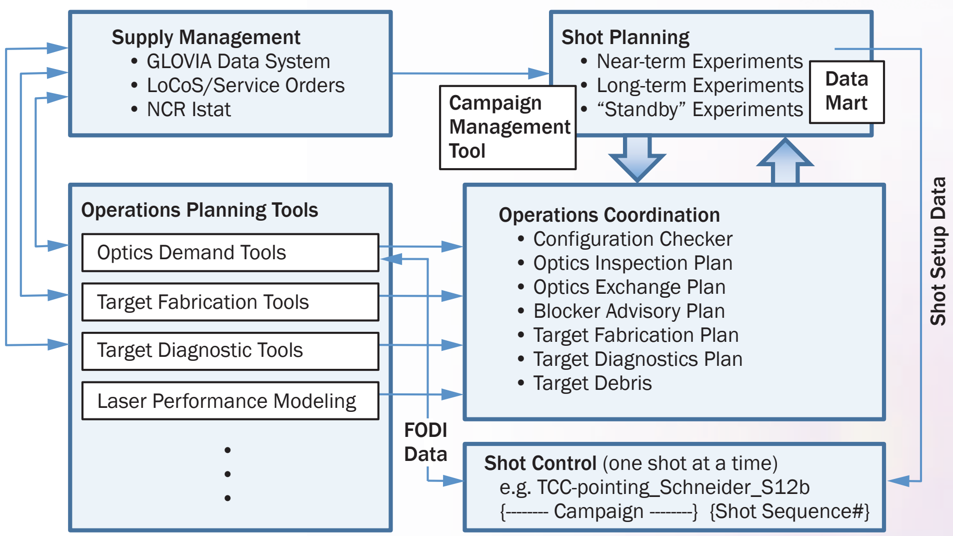

High-level flow diagram depicting SPLAT's integrated shot planning, operations, and primary supply and demand loops. Operational planning will allow NIF to forecast demand for optics, targets, and target diagnostics.

The optics demand component is built using models that predict the damage that might occur to the final optics from the proposed shots. The predictive models integrate calculated energetics, current optics status, and physics-based mathematical damage "rules" that predict optics flaw growth and damage initiation. These models are then run on a multinode analytical cluster, and the results are used to produce decision-support reports on optics inspection planning, optics exchanges, and optics damage blocker placement advisories.

The SPLAT Project will eventually provide a multitude of decision-support and operation optimization tools. To maintain their aggressive delivery schedule and respond to evolving requirements, the SPLAT team has combined agile-based project planning and management tools with a highly modular system infrastructure and component architecture. The team works closely with the project stakeholders in short, iterative cycles to deliver the requested features and capabilities on time.

RETURN TO TABLE OF CONTENTS

CASTLE Improves Nuclear Weapon Safety-Basis Processes

Contact

Mark Piscotty,
piscotty3@llnl.gov

Valerie Noble,
noble22@llnl.gov

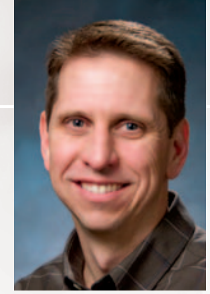

Piscotty

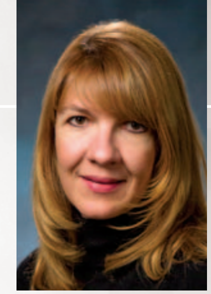

Noble

Summary

An essential part of NNSA's nuclear stockpile stewardship responsibility involves safely conducting weapons operations, such as disassembly and inspection, at the Pantex Plant in Texas. The primary objective of the Collaborative Authorization for the Safety-Basis Total Lifecycle Environment (CASTLE) Project is to transform cross-site, safety-basis development and management processes in support of these activities. LLNL leads the multisite, multiyear project and works with stakeholders from Pantex and Sandia and Los Alamos national laboratories to develop and incorporate process improvements and standardizations into technical solutions that assimilate weapons operations information with a set of safety-basis controls and weapons response data. In 2009, the CASTLE team developed an innovative software application that allows engineers to identify risks early in the development process, facilitates the design of tooling and process flow, and enables the creation of safety-basis evidence that meets DOE standards and directives.

Progress in 2009

Safety-basis assembly/disassembly and inspection processes at the Pantex Plant integrate operating procedures with hazard identification, weapons response, and accident analyses to form a safe, efficient, and effective operating environment. Many safety-basis processes were rooted in manual, labor-intensive data operations until 2006, when a series of value stream analyses concluded that improving the information technology infrastructure would significantly increase safety-basis efficiencies. With this goal in mind, the CASTLE Project was initiated in April 2007. Its products leverage a limited prototype version (CASTLE-Lite), which was successfully used for the W80 Seamless Safety in the 21st Century (SS-21) program. CASTLE is composed of two, tightly coupled products: CASTLE-PX, a Pantex-centric product focused on improving process flow and authorization basis, and CASTLE-WR, a design agency-centric product focused on generating weapons response.

CASTLE-PX provides a single, Web-based platform that combines weapons operations information (process flow) with the safety-basis control set and weapons response. CASTLE-PX allows process designers to identify risks early in the safety-basis development process, and it facilitates the process flow definition and hazard identification, characterization, and mitigation. Once the process data is defined and validated, CASTLE-PX facilitates the generation of the hazard analysis report (HAR). Although CASTLE-PX is not scheduled for full production deployment until early 2010, preproduction releases were successfully deployed throughout 2009. CASTLE-PX is now being used to support W84, B53, W76, and B83, and it is providing marked improvement in both the data integrity and management of Pantex's safety basis.

Like CASTLE-PX, CASTLE-WR provides a single Web-based platform with multisite access. It is designed not only to integrate multisite weapons response but to enforce a standardized language and common methodologies, which ultimately lead to more consistent weapons response across weapons systems. CASTLE-WR is still in the early stages of development, but similar to CASTLE-PX, the team will deliver incremental phased releases throughout the development lifecycle.

In September 2009, the CASTLE Project completed a Level 2 milestone by confirming the delivery of critical software functionality within the CASTLE-PX product. This milestone represents the culmination of a remarkable team effort forged with software development skill, extensive domain knowledge, embedded customer support, and an unwavering drive to meet customer and NNSA needs.

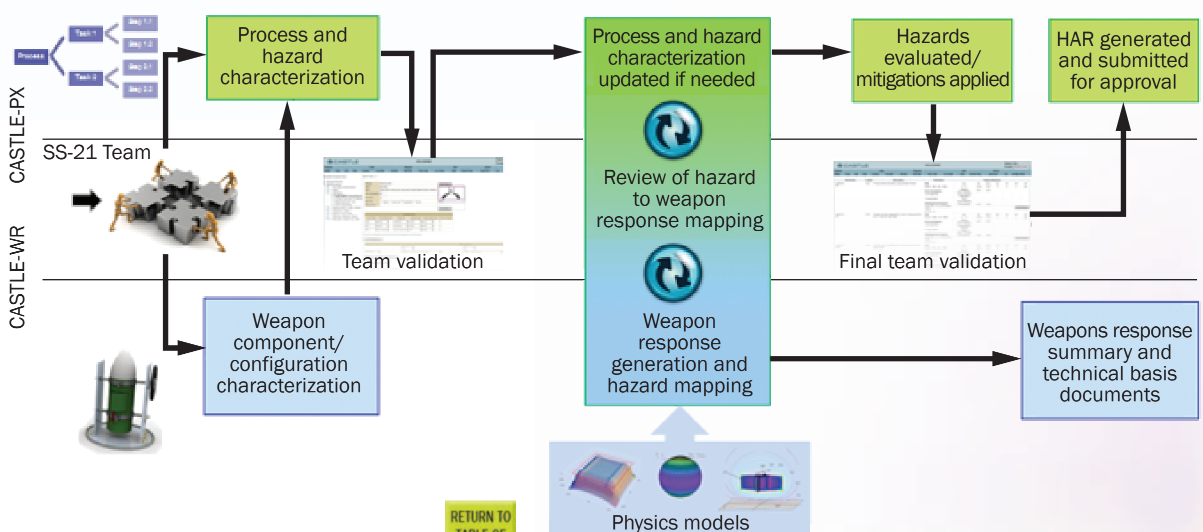

CASTLE products facilitate collaborative safety-basis development through data management and engineering analysis capabilities.

RETURN TO TABLE OF CONTENTS

Information Technology and Security

Computation has provided our large and diverse laboratory with information technology (IT) support for many years. Computation people develop strategies, lead projects, and direct operations in cyber security, desktop/server support, networking, customer support, and communication and collaboration services. Much of the Laboratory's critical IT infrastructure is entrusted to Computation personnel, as is oversight of the Cyber Security Program, which ensures the security of Lawrence Livermore National Laboratory's (LLNL) unclassified and classified IT assets. This section includes highlights from our IT accomplishments in 2009. The articles showcase our continuing success in helping the Laboratory negotiate the use of modern, secure, and cost-efficient IT capabilities.

Computation has a long history of supplying people with IT skills to individual programs, but more frequently, we are delivering IT solutions that serve the entire Laboratory. Replacing many small, competing services with one site-wide service can help users be more productive, reduce costs, and improve security. This year Computation personnel completed a project to automate desktop and server management. System administrators can now do more in less time and in a more secure manner. We also originated the idea for and enacted a centralized, single-call help function, which addresses a spectrum of issues ranging from Microsoft Office questions to network problems. These two successful efforts paved the way for other site-wide service deployments, one of which is the Blue Network. The Blue Network is designed to provide secure access to network resources for non-U.S. citizen employees and collaborators. This service has been deployed to more than 25 buildings at LLNL and has improved security and service for its users.

Computation, in collaboration with the Chief Information Office (CIO) Program, has advanced a model of unified communications to help Laboratory staff partner with others in an increasingly fast-paced world. The model calls for integrating real-time tools—voice, data, and video—into a single, intuitive user interface and has proven to help companies work more efficiently and save money. Unified communications solutions help people find, reach, and collaborate with others. LLNL is in the midst of a multiyear effort to provide users the conveniences of modern collaboration tools: e-mail, instant messaging, and smartphones (e.g., BlackBerrys and iPhones). This deployment is a milestone for LLNL; it is the first time employees will have access to managed smartphone service. However, more work is needed to prepare LLNL for unified communications. Next year, we hope to report further progress toward this goal.

LLNL does its best work when integrated project teams come together to solve tough problems. In this regard, IT is no different. One of Computation's focuses this year was to build more effective teams of IT and cyber security personnel. This section includes three articles that exemplify the success of our efforts. In one example, an integrated team engineered an approach to automate the management of security configurations. The idea was first developed by IT personnel, and then Cyber Security staff were brought in to help turn it into a production service. The collaborative nature of the project helped ensure that it would receive widespread support from both IT and Cyber Security management. Another example of an integrated project team's success is the BlueCoat Blocking tool, which protects the Laboratory's unclassified network by blocking vulnerable systems from the Internet. In another project, the Cyber Security and IT organizations worked together to help mobile users better protect their data. The danger of data loss from LLNL's increasingly mobile staff was becoming a risk too great for the institution to accept. In response, Computation quickly developed two solutions: the Laptops on Foreign Travel service and disk encryption for Windows systems. The close collaboration between Cyber Security and IT helped ensure that both security and usability issues were considered in the design of the new services.

Computation also lead the Laboratory's first-ever automated deployment of a major application, Microsoft Office 2007/2008. After months of preparation, the deployment to 8,500 desktops was carried out over a single weekend. The successful completion of this project required disciplined planning and execution within the IT organization and effective communication to organizations across the Laboratory.

As IT services evolve at the Laboratory, Computation will continue to anticipate, innovate, and deliver the tools and infrastructure necessary for employees to accomplish their mission goals.

TERRI QUINN,
QUINN1@LLNL.GOV

RETURN TO TABLE OF CONTENTS

Security Services Improve LLNL Data Protection

Contact

Robyne Teslich,
teslich2@llnl.gov

Summary

In 2009, LLNL's CIO Program continued to enhance security services to protect mobile devices and mobile data. Mobile devices present a greater risk of data loss due to device loss or theft. To reduce the risk of data loss and address requirements in the National Nuclear Security Administration (NNSA) policies, the CIO Program implemented full-disk encryption on mobile devices and removable media.

A broader set of threats exists when a mobile device user travels to a foreign country, where there may be more restrictions on encryption or more risk from foreign adversaries. To address the additional risks during foreign travel, the CIO Program launched an institutional service, called Laptops on Foreign Travel (LOFT), which is administered by the Computation Directorate.

Progress in 2009

Recent headlines about major companies losing sensitive data illuminated the need to protect LLNL data when it leaves the Livermore site. Barring the elimination of mobile devices and laptops leaving the site altogether, the best way to protect data on mobile devices is to implement full-disk encryption. File- and folder-level encryption rely too heavily on the user knowing where sensitive data is located on the device at all times.

In spring 2009, a multiphase project was initiated to implement encryption on laptops, Universal Serial Bus (USB) thumb drives, removable media, and personal digital assistants. The majority of LLNL laptops are Windows laptops, so they were addressed first. By the end of 2009, the GuardianEdge full-disk encryption product was installed on more than 90% of the targeted Windows laptops. Macintosh OSX and Linux encryption solutions will be addressed during the next phase of the project, followed by USB thumb drives and other removable media.

To address the additional risks facing laptops that enter foreign countries, LLNL followed the path of several other laboratories and implemented a centralized pool of laptops for users to take on travel. The LOFT program mitigates the risk of systems entering foreign countries by securely configuring the laptop prior to travel, minimizing laptop access to LLNL networks, scanning systems before and after each trip, and minimizing the amount of data a user takes on travel. The program has been well received by foreign travelers and used heavily since its inception in October 2009. Both Windows and Macintosh OSX platforms are available for users to take on their trips. Frequent travelers or those that require unique laptop configurations have the option to pay for a dedicated LOFT laptop. The dedicated laptop is maintained by the LOFT pool administrator but is always available for use only by the owner.

The security services implemented in 2009 have significantly reduced LLNL's risk of data loss with little impact to user productivity. Additional services will be deployed in 2010 to address more platforms and ensure that all sensitive LLNL data that leaves the site is protected.

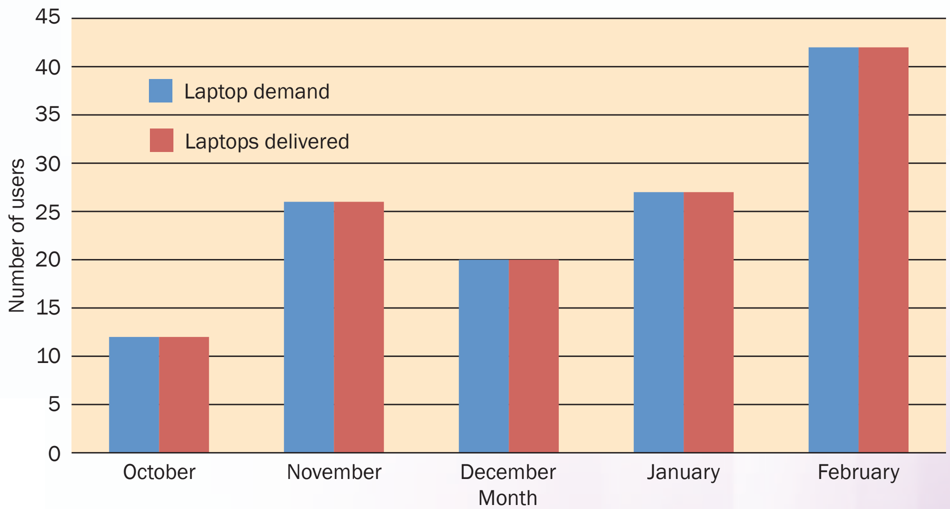

LOFT program usage since its inception in October 2009 through February 2010.

RETURN TO TABLE OF CONTENTS

Contact

JEAN SHULER,
SHULER1@LLNL.GOV

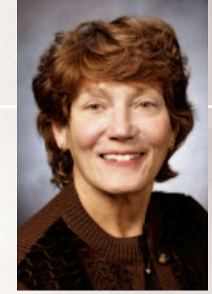

Automated Tools Enable Simultaneous Upgrade to Office 2007/2008

Summary

The Microsoft Office 2007/2008 suite was successfully deployed to almost all desktop machines at the Laboratory over a three-day weekend in January 2009. The deployment was the first time a major software application was distributed and installed simultaneously across the enterprise using automated tools. The Office 2007/2008 suite was distributed to 8,500 of LLNL's 9,800 desktop (Windows and Macintosh) machines using the institution's automated tool, LANDesk. Installation on the remaining machines was complete in February.

The upgrade to Office 2007/2008 allows LLNL scientists to more efficiently collaborate with government agencies and external partners who already use the suite of applications. Since the deployment, the Laboratory's IT centralized support services are better able to rapidly patch systems and reduce vulnerabilities. The migration also helped prepare the Laboratory for its next large software deployment scheduled for completion in 2010—Microsoft Exchange.

Progress in 2009

Deploying Office 2007/2008 was imperative to keep the Laboratory current with standard technologies and to increase collaboration opportunities and user productivity. Formal planning began in March 2008 when a project team was established to discuss and minimize risks, set milestones, and assure deliverables were met. The team decided to upgrade users at the same time to mitigate document interchange problems. According to industry best practices, a simultaneous migration is more effective than a gradual, optional change because the file formats for Office 2007/2008 are markedly different than previous versions.

Once the CIO approved the project proposal in August 2008, an aggressive six-month timeline commenced. The project team surveyed users to determine their readiness for the deployment and wrote scripts and software tools to gather and analyze information. They also implemented a Laboratory-wide communication plan, which included a project Web site, feedback mechanism, and frequent articles in LLNL's e-newsletter *NewsOnLine*. The team offered extensive in-person, televised, and online training resources, both for technical staff and for mainstream users.

One of the critical aspects of the migration planning was an early adopters program, in which dozens of users tested the migration process on multiple platforms and provided considerable and candid feedback to help refine the distribution, service, and support methods. One week before the deployment, project leaders met to decide whether or not to proceed with the upgrade, and their decision to proceed was unanimous. IT staff and technical experts from each of LLNL's five principal directorates set up "war rooms" to troubleshoot issues during the weekend push of the software and Computation's Institutional Service Desk (4HELP) personnel were trained and prepared for an influx of customer inquiries after the weekend migration.

In the year since the deployment, users have become more familiar and comfortable with the new, substantially different features in Office 2007/2008. The successful deployment experience is one of many examples of how Computation delivers on its commitment to provide stable, quality IT products with an integrated solution to LLNL users.

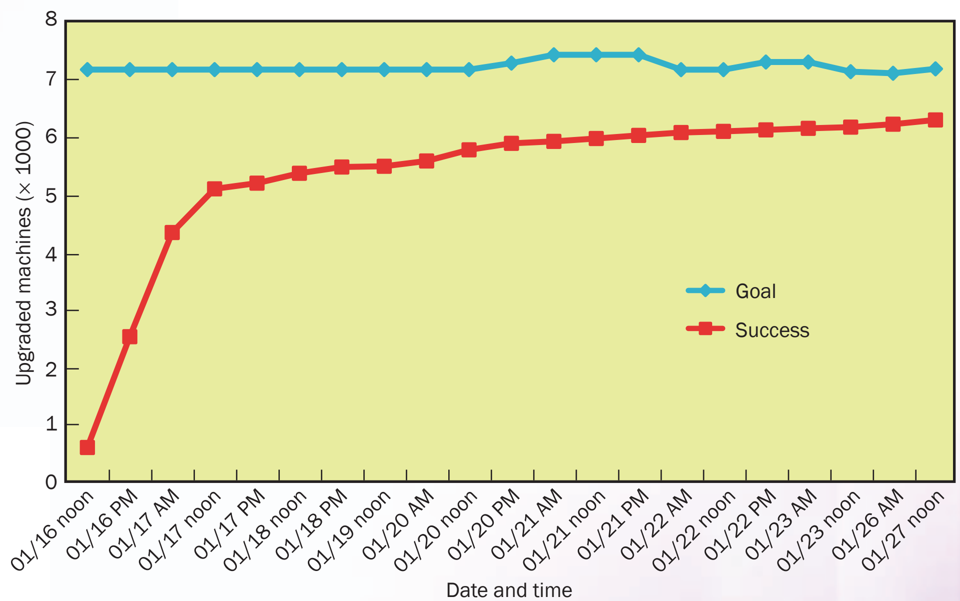

LLNL's LANDesk automated tool was key to the Microsoft Office 2007/2008 upgrade success. More than 75% of the Office clients were automatically installed using LANDesk in the first 24 hours.

RETURN TO TABLE OF CONTENTS

LLNL Expands Collaboration Opportunities through IT Services

Summary

Modern collaboration tools are becoming less of a convenience and more of a necessity in today's research and development environments. Scientists, researchers, and support staff need tools that manage and organize e-mail, calendars, and tasks through an integrated interface, and they need the ability to communicate quickly through instant messaging (IM). The LLNL workforce is increasingly mobile, and as such, they require access to the full suite of adaptable, real-time communication and collaboration capabilities, including smartphones, desktop clients, and Web interfaces.

In 2009, the CIO Program and Computation Directorate made headway on three major projects to improve LLNL's collaboration infrastructure: replace the existing, outdated e-mail and calendar solutions with Microsoft Exchange; provide BlackBerrys for mobile users; and offer an enterprise IM service. The Exchange migration project is 40% complete and will culminate with a weekend upgrade of the calendar system in summer 2010. BlackBerrys will be available to Exchange users beginning in spring 2010. The IM service (Office Communication Server) was deployed in November 2009 and is used daily by more than 3,500 people.

Progress in 2009

LLNL's current e-mail and calendaring tools are out of date—the vendor no longer supports them, and they do not provide the functionality users have come to expect in a modern communication environment. In addition, the current e-mail system does not integrate with commercially available electronic records retention solutions, making it impossible for the Laboratory to abide by electronic records retention requirements.

After analyzing LLNL requirements and the collaboration tools marketplace, Microsoft Exchange was selected as LLNL's new e-mail and calendar solution. Exchange is a leader among its competitors, and it provides a feature-rich environment in which users can communicate and collaborate. As an industry leader, Exchange opens the door to a plethora of third-party solutions for electronic records retention systems, voicemail, and other collaboration capabilities.

The Exchange project began in 2008 and was fully funded in October 2009 to support the level of service required by the Laboratory's principal directorates. The system provides an account availability of 99.9% and gives each user 1 gigabyte (GB) of mailbox storage, which doubles the server-side storage per user. Premium storage (6 GB) is available to approximately 25% of the e-mail user population but requires approval by a principal directorate's IT manager.

Migrating 7,800 accounts to a new e-mail and calendar system, while ensuring ongoing user productivity and no loss of data, is an intricate endeavor that requires significant personnel effort. The Laboratory has a complex IT ecosystem, consisting of Windows, Macintosh, and Linux desktop and laptop configurations scattered across its organizations. One of the early challenges we had to overcome was identifying solutions that met the needs of the various users. The selected e-mail client solutions are Apple Mail for Macintosh users, Outlook for Windows users, and several options for Linux users:

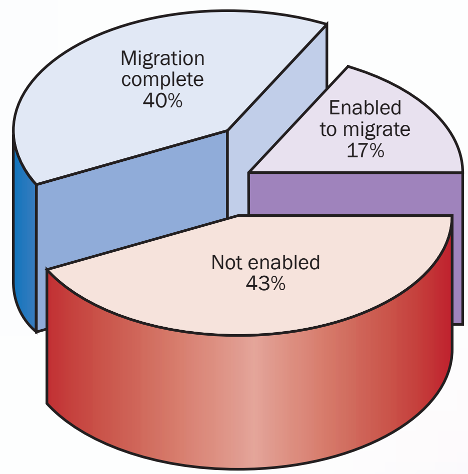

Forty percent of LLNL users migrated to Exchange during the first three months of the process. The remaining users will migrate by mid-2010.

Outlook using CrossOver or another Windows virtual environment, Outlook Web Access, or Internet Message Access Protocol (IMAP).

LLNL users often encrypt their e-mails using Entrust, which added another layer of complexity to the migration process. Before the migration could begin, Windows and Macintosh systems needed new versions of Entrust installed. In addition, we collaborated with Entrust to develop a special utility to help users migrate old Entrust-encrypted e-mails stored on their desktops to the new client.

We developed extensive documentation to help end-users navigate the migration process and new tools and trained technical support staff to minimize the impact of the change on user productivity. We also created tools that allow users to control the timing of their e-mail migration, including the migration of the e-mail stored on their desktop. Windows users began migrating to Exchange in December 2009, and Macintosh and Linux users will begin migrating in March 2010.

All users must migrate to Exchange before the calendar migration can begin. The interconnected nature of calendaring requires that users migrate at the same time. This migration is scheduled for a weekend in early summer 2010.

LLNL's production BlackBerry service will be available in April 2010. For the first

Contact

REA SIMPSON,
SIMPSON6@LLNL.GOV

time, users will have access to e-mail, IM, and the Web from anywhere via the cellular network. Calendars will be integrated with the BlackBerry service after the institutional calendar migration in summer 2010. LLNL-managed BlackBerrys will be configured securely to ensure LLNL data is protected.

In October 2009, the CIO launched an institutional IM service based on the Microsoft Office Communication Server product. Prior to this service, isolated IM servers were being used by some LLNL organizations and external IM services (e.g., AOL and Yahoo) were blocked due to security concerns. The new IM solution is integrated with the Outlook e-mail client so users can quickly and easily respond to an e-mail using IM. A “presence awareness” function allows users to easily check whether or not a coworker is available for an IM session.

Employees have come to expect a more advanced set of communication and collaboration tools in recent years. The CIO Program and Computation are responding by adding robust collaboration tools that substantially improve the Laboratory’s base e-mail and calendaring solutions and by offering additional services, such as BlackBerrys and integrated IM. These improvements will significantly increase user productivity and provide a foundation for future collaboration enhancements.

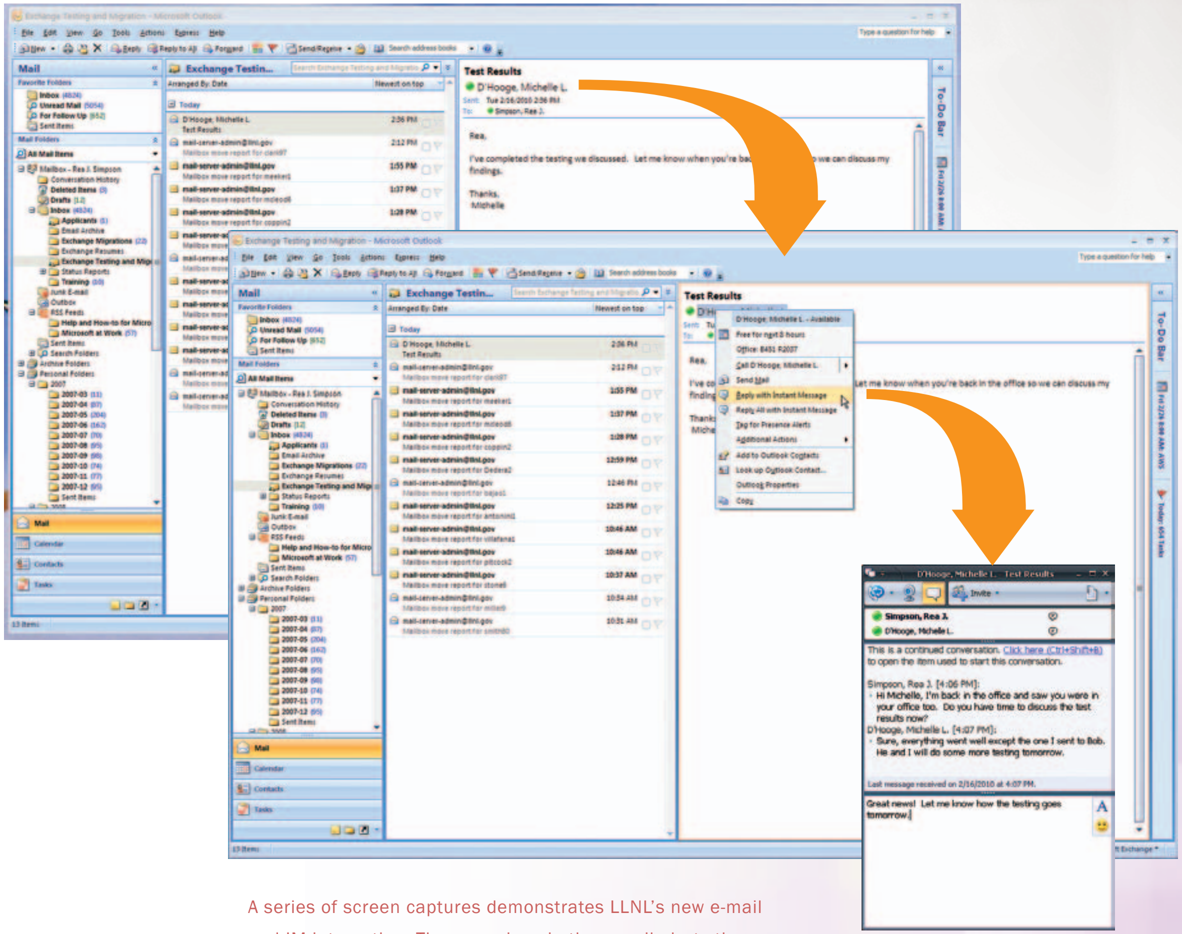

A series of screen captures demonstrates LLNL’s new e-mail and IM integration. The green icon in the e-mail alerts the recipient that the sender is available via IM. A menu pull-down in the e-mail initiates the IM session so the conversation can continue in real time.

RETURN TO TABLE OF CONTENTS

Blue Network Simplifies Network Access Management

Contact

ROBIN GOLDSTONE,
GOLDSTONE1@LLNL.GOV

ENS.LLNL.GOV/BLUENETWORK

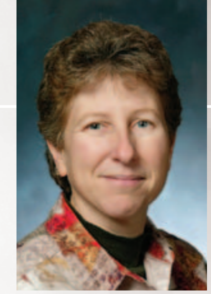

Summary

The Blue Network was deployed in 2009, marking the culmination of a multiyear project to provide a secure network access model for LLNL Foreign National (FN) employees and collaborators. The Blue Network is both a technical solution and a functional framework that manages the cyber access needs of FNs. More than 200 individualized computer security (CyberTrak) plans have been closed since May 2009 as a result of the Blue Network deployment. The Blue Network delivers an enhanced security posture and substantial cost savings for LLNL and improves the productivity of the FN workforce.

Progress in 2009

The Blue Network consists of a network segmentation architecture utilizing Virtual Local Area Network (VLAN) technology and a highly scalable role-based firewall that enforces network access policy between Blue VLANs and other network segments (e.g., the "Yellow" Restricted Open Network). Remote access to the Blue Network environment is provided through a new virtual private network access group (VPN-B). The technical solution was developed and prototyped in 2007 and was implemented and piloted in 2008.

In December 2008, LLNL received approval from NNSA for the Blue Network to begin production.

Production deployment could not commence until the relevant policies and procedures were reviewed and aligned with the Blue Network model. This process required more than simply revising a few documents; it meant rethinking the entire way that FN cyber access would be managed. One outcome of these deliberations was the decision to use LLNL's Configuration Management Database (CMDB) to track systems approved for FN access, a process formerly managed by the CyberTrak system. In less than a month, CMDB was upgraded to provide the necessary functionality. By the end of March 2009, most of the policy revisions were complete and production could proceed.

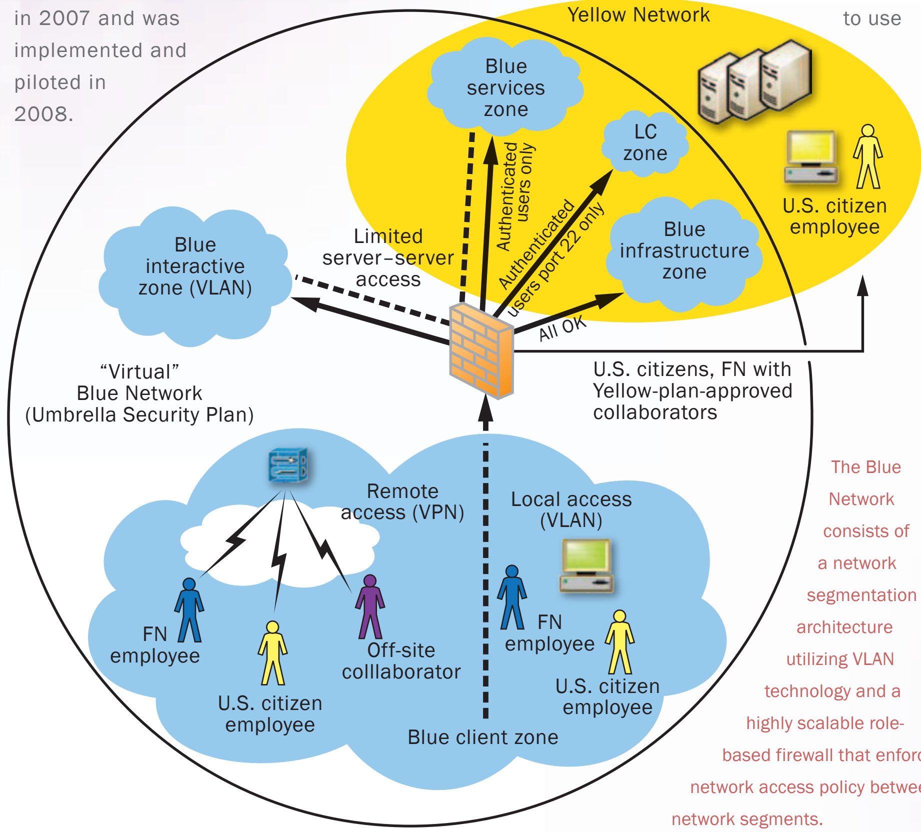

The Blue Network consists of a network segmentation architecture utilizing VLAN technology and a highly scalable role-based firewall that enforces network access policy between network segments.

The Science and Technology (S&T) Principal Directorate, which has the largest number of FNs at the Laboratory, became the first programmatic organization to use the Blue Network. We established a project timeline and priorities and deployed the service on a per-building basis. The multistep deployment process included reconfiguring 25 building network infrastructures to support Blue Network VLANs, deploying the Blue Network firewall, identifying and reconfiguring systems to be added to the Blue Network, and validating that FN users could access the necessary cyber resources. Once it was established that the users had the same access they had on their CyberTrak plan, their plan was submitted to LLNL's Computer Security Program for closure.

While the Blue Network was being deployed to FN employees in the S&T Principal Directorate, we were transitioning off-site FN collaborators to the new VPN-B service. As of December 2009, the Blue Network has been deployed to 25 S&T office buildings, server rooms, and laboratories, resulting in the elimination of more than 200 CyberTrak plans and saving an estimated $2 million annually. The Blue Network will be deployed to other programmatic customers in early 2010.

RETURN TO TABLE OF CONTENTS

BlueCoat Blocking Improves Laboratory Cyber Security Posture

CONTACT

BRETT CORNELL,
BCORNELL@LLNL.GOV

SUMMARY

LLNL must maintain a secure computing environment to accomplish its mission. To this end, the Laboratory has implemented many cyber defenses over the years, including intrusion detection systems, proxy servers, firewalls, and systems management tools. This year, LLNL continued to improve its cyber defenses by introducing a new system that permits only security-compliant computers to access the Internet. Systems not in compliance with LLNL cyber security policies cannot access the Internet until their security deficiencies are corrected. By denying noncompliant systems access to the Internet, LLNL has closed a significant avenue of cyber attacks and improved its overall security posture.

PROGRESS IN 2009

In summer 2009, LLNL introduced a new system to its cyber security arsenal, the BlueCoat Blocking system, or BCB. BCB gathers and processes information from a variety of sources—Active Directory, the Laboratory's Configuration Management Database, and LANDesk infrastructure, among others—to ascertain a computer system's cyber security profile. BCB issues "soft blocks" on systems it determines are not in compliance with LLNL's cyber security policy. Systems on which soft blocks have been issued cannot access the Internet. They can, however, access resources on LLNL's internal computer network. If a blocked system attempts to use the Internet, a Web page appears telling the user that the system is blocked for failure to comply with cyber security regulations and enumerates why the system is blocked (e.g., LANDesk is not installed or is not reporting properly, the system is missing required security patches, its security scans are outdated). The soft block remains in effect until the system has been brought into compliance with all security requirements. Once BCB determines that the system is compliant, it automatically removes the soft block.

Each principal directorate's Organizational Information Systems Security Officer (OISSO) can grant a system a "blocking exemption" that will prevent it from being blocked for a defined period of time. The OISSO can only grant exemptions to systems under his or her authority. A system with an exemption remains unblocked until the exemption expires.

BCB provides system administrators with a set of screens that display systems that are currently blocked (the Master Block List) and those that will be blocked (the Pending List). Basic information regarding a blocked system is also displayed (e.g., Internet Protocol address, Domain Name System name, primary user). BCB provides search capabilities to help system administrators identify systems that are blocked or will soon be blocked. In addition, system administrators can run reports against the underlying BCB database to provide additional information to users. The system is accessible through the Web only to authorized users.

Soon after the successful introduction of BCB, Laboratory management expressed interest in increasing its capabilities. During the remainder of 2009, functional requirements for a new version of BCB were developed. The new version of the system will include more sophisticated algorithms for determining which systems should and should not be blocked, and it will enforce additional security requirements on Macintosh and Linux systems. It will also have the ability to "hard block" systems, which will completely prevent them from using LLNL's network. The new version of BCB will be deployed in 2010.

The home screen for BCB, one of LLNL's cyber security applications. From the main page, system administrators can run various reports and search for system compliance.

RETURN TO TABLE OF CONTENTS

Contact

Sue Smith,
smith74@llnl.gov

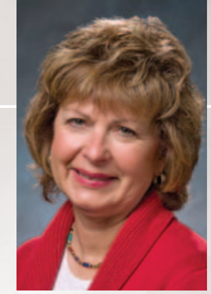

Centralized Certification Library Fulfills Compliance Requirement

Summary

Computation has made important contributions to LLNL's NNSA Policy Letters (NAPs) compliance efforts. NAPs are a set of orders that enact prescriptive controls over how NNSA organizations implement and manage their cyber security defenses. In partnership with Cyber Security and Laboratory programs, Computation launched a centralized library of certified configurations for classified and unclassified systems, called the Site Security Component Library (SSCL). SSCL brings the Laboratory into compliance with NAP 14.1-C, Appendix B and is expected to save the Laboratory time, money, and effort. Most importantly, it frees precious resources that can be redirected toward meeting the institution's national security missions.

Progress in 2009

SSCL was deployed to several organizations at the Laboratory in 2009. SSCL is a collection of reusable certified configurations that can be inherited by unclassified and classified system security plans. For example, if a particular version of an operating system configuration resides in SSCL, it can be quickly and easily installed on a set of desktop computers. Each component that resides in the library consists of several parts, including the configuration itself, the authoritative basis for the configuration, the NAP controls, and the testing tool.

The library, which was built in accordance with software engineering and management best practices, is managed by a clear separation of roles and duties and can be audited. SSCL has documented processes for lifecycle management, significant security changes, and inheritance. Significant security changes require additional reviews and are subject to a formal "acceptance of risk" by NNSA's Livermore Site Office (LSO).

A component can be submitted to the library for consideration by following a documented process. The SSCL manager (SSCM) is key to the success of the library. The SSCM works with the component originator and ensures the component package is complete. The SSCM can use the testing tool to verify that the configuration test is repeatable and that it properly interrogates the configuration. Once the SSCM is satisfied with the component package, the package is given to the Cyber Security Site Manager (CSSM).

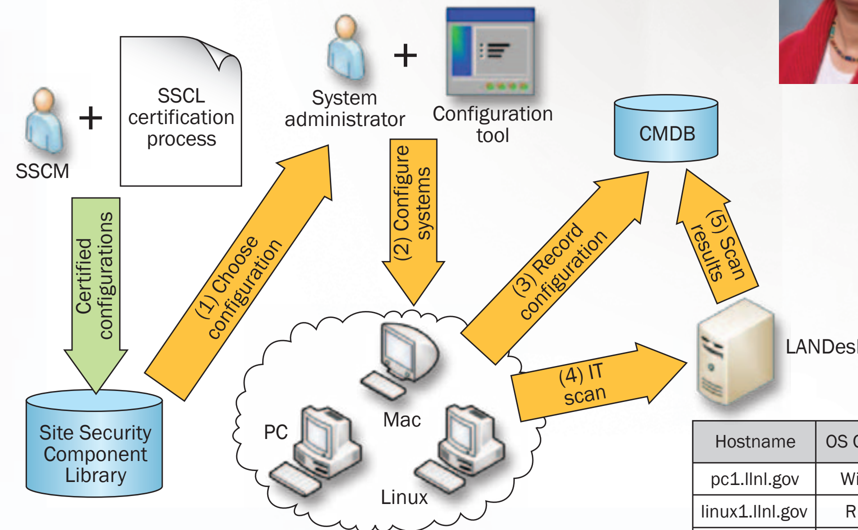

Hostname	OS Configuration
pc1.llnl.gov	WinXP_FDCC
linux1.llnl.gov	RHEL5_CIS
mac1.llnl.gov	MacOSX10.5_CIS

Configuration Name	Configuration Tool	Configuration Guide
WinXP_FDCC	fdcc-winxp-xccdf.xml	FDCC_XP_Guide.pdf
RHEL5_CIS	llnl-fortify-rhel5.conf	RHEL5_NIST_Guide.pdf
MacOSX10.5_CIS	fortifyX-macosx-10.5.conf	CIS_OSX_10.5_Benchmark_V1.0.pdf

The process of applying a configuration from SSCL begins when the system administrator chooses a configuration from the SSCL database using the configuration tool. The configuration is then applied to a set of machines and stored in CMDB. The LANDesk software interrogates the system(s) during an automated scan and records the configuration it finds on the system(s). The results of the scan are then recorded in CMDB for future validation and verification.

The CSSM reviews the package for security significance and residual risk. If the package approaches or exceeds the risk boundary of LLNL's current certification, the package is submitted to LSO for assumption of risk. Packages that are deemed by LSO as "security significant" are well documented and described so processes are repeatable and not prone to misinterpretation. Once a component is certified and accredited, the component is "checked in" to the library and available to use in any information system security plan.

The SSCM monitors and publishes metrics regarding the lifecycle of each component. The library automatically "checks out" components that are no longer in a supported lifecycle. Once a component has reached the end of its lifecycle, it is the responsibility of the security plan manager to bring the component back into compliance.

SSCL provides a necessary, standardized, and automatic process for deploying reusable IT configurations throughout the Laboratory in a way that helps the institution meet its cyber security obligations. As the library of accepted components grows, more customers will see their cyber security needs met by SSCL.

RETURN TO TABLE OF CONTENTS

Appendices

RETURN TO TABLE OF CONTENTS

Industrial Collaborators

Company	Topic	LLNL contact
AbsInt	Worst-case execution time analysis	Dan Quinlan
Adaptive Computing Enterprises, Inc.	Moab Workload Manager	Don Lipari
Allinea Software	Scalable debugging infrastructure	Greg Lee
Appro	Scalable capacity clusters	Mark Seager
Aramco	Oil well reservoir visualization	Eric Brugger
Arista Networks	Low-latency Ethernet networks	Matt Leininger
Bosch	Worst-case execution time analysis	Dan Quinlan
Bull, Commissariat à l'Énergie Atomique, Hewlett-Packard	SLURM resource management software	Morris Jette
Cisco Systems	Application engine in router	Maya Gokhale
Cisco Systems, Dell, DataDirect Networks, Intel, LSI Industries, Mellanox, Qlogic, Red Hat, Sun Microsystems, Supermicro	Hyperion collaboration	Mark Seager
Commissariat à l'Énergie Atomique, NNSA (LLNL, SNL, LANL)	Deploying high-performance computing to produce more precise and reliable simulations	Kim Cupps
Convey Computer	High-performance reconfigurable computing	Maya Gokhale
Cray	Programming environments	Bronis de Supinski
Cray	Scalable systems, open-source strategy	Mark Seager
DataDirect Networks	RAID 6 R&D for I/O systems	Brent Gorda

Industrial Collaborators (continued)

Company	Topic	LLNL contact
DataDirect Networks	RAID application engines	Maya Gokhale
IBM	High-performance storage system	Jerry Shoopman
IBM	Infosphere streams	Maya Gokhale
IBM	Scalable systems, multiple areas	Mark Seager and others
IBM	Seismic modeling	Shawn Larsen
Imigene	DNA signatures for blood-borne diseases	Tom Slezak
Impulse Accelerated Technologies	C-to-FPGA compiler	Maya Gokhale
Instrumental Corp.	Storage systems	Maya Gokhale
Intel Corporation	Optimization of seismic algorithms	Shawn Larsen
Intel Corporation	Solid-state memory	Maya Gokhale
Kaman Precision Products	Yield selection of Mk82	William Bateson
Microfluidic Systems	Pathogen signatures	Tom Slezak
Nallatech	Reconfigurable computing board	Maya Gokhale
Navistar Corp., Michelin	Heavy vehicle aerodynamic drag reduction	Kambiz Salari
NVIDIA	3D-mesh compression	Jon Cohen
OpenFabrics Alliance, Mellanox, QLogic	OpenFabrics Enterprise Distribution	Ira Weiny
OpenMP Consortium	Shared-memory programming models	Bronis de Supinski

Industrial Collaborators (continued)

Company	Topic	LLNL contact
OpenWorks	Valgrind memory tool and threading tool development	John Gyllenhaal
Red Hat	Operating systems	Mark Seager
Samsung	Solid-state memory	Maya Gokhale
Schlumberger	Synthetic seismic data set generation	Shawn Larsen
Siemens	Wind power	Bill Henshaw
SRI International	Acoustic and meteorological modeling	Shawn Larsen
Sun Microsystems	Lustre file system development and deployment	Brent Gorda and Mark Gary
Sun Microsystems	Seismic modeling on grid systems	Shawn Larsen
Terascala	High-performance I/O subsystems	Brent Gorda
TotalView Technologies	Parallel debugger scalability and enhanced memory tools	Dong Ahn and Scott Futral
URS Corporation	Earthquake hazard estimation	Shawn Larsen
Virident	Storage class memory	Maya Gokhale
XtremeData, Inc.	Field-programmable gate array accelerator, database engine	Maya Gokhale

RETURN TO TABLE OF CONTENTS

Academic Outreach

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
UC Berkeley	James Demmel	Joint Research	Towards optimal petascale simulations	ASCR SciDAC	Robert Falgout
UC Berkeley	Doug Dreger	Subcontract	Earthquake hazard	IGPP	Shawn Larsen
UC Berkeley	Kurt Miller and Michael Jordan	Collaboration, Lawrence Scholar Program	Latent variable models	LDRD	Tina Eliassi-Rad
UC Davis	Francois Gygi	Subcontract	Algorithms for electronic structure and first-principles molecular dynamics simulations using large-scale parallel computers	ASC	Art Mirin
UC Davis	Zhendong Su	Subcontract	Software security analysis research	LDRD	Dan Quinlan
UC Davis	Ken Joy	Joint Research	Visualization and Analytics Center for Enabling Technologies	ASCR SciDAC	Mark Duchaineau
UC Davis	Ken Joy	Subcontract	Semantic analysis of large-scale aerial video	LDRD	Mark Duchaineau
UC Davis	Ken Joy	Collaboration, Lawrence Scholar Program	Discrete multimaterial interface reconstruction for volume fraction data	UCOP	Mark Duchaineau
UC Davis	Ken Joy	Collaboration, Lawrence Scholar Program	Video processing research for the VidCharts and Persistics projects	UCOP	Mark Duchaineau
UC Davis	Nelson Max	Joint Research	3D from wide-area aerial video	DOE Nonproliferation	Mark Duchaineau
UC Davis	Soheil Ghiasi	Collaboration	Network security	LDRD	Maya Gokhale
UC Davis	Kwan-Liu Ma	Joint Research	Large-graph data visualization	LDRD	Peer-Timo Bremer
UC Davis	Bernd Hamann	Joint Research	Analysis and visualization of scientific data using topology-based methods	LDRD	Peer-Timo Bremer
UC Davis	Kwan-Liu Ma	Joint Research	Visualization of RADAR tomography reconstruction	DoD	Mark Duchaineau
UC Los Angeles	Alfonso Cardenas	Collaboration, Lawrence Scholar Program	iScore: measuring the interestingness of articles in a limited user environment	UCOP	David Buttler
UC Los Angeles	Thomas Grismayer	Collaboration	Comparison of particle and continuum kinetic simulations of laser–plasma interaction		Jeff Hittinger

Academic Outreach (continued)

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
UC Los Angeles and Weizmann Institute (Israel)	Achiezer Brandt	Consultant	Geometric and algebraic multigrid techniques	ASCR Base	Robert Falgout
UC Merced	Qinghua Guo	Joint Research	Learning from presence-only data	UCOP	Tina Eliassi-Rad
UC Riverside	Michalis Faloutsos	Collaboration	Behavioral approaches to traffic classification	LDRD	Tina Eliassi-Rad
UC Riverside	Marios Iliofotou	Collaboration	Behavioral approaches to traffic classification	LDRD	Tina Eliassi-Rad
UC San Diego	Allan Snavely and Laura Carrington	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
UC San Diego	Scott Baden	Subcontract	Data-driven execution of latency-tolerant algorithms	LLNL Overhead	Dan Quinlan
UC San Diego	Tim Barnett, Michael Norman, and Randy Bank	Subcontract	LLNL–UCSD scientific data management	UC Management Contract	James McGraw
UC San Diego	Falko Kuester	Joint Research	Large-scale atomistic simulation visualization	ASC	Mark Duchaineau
UC San Diego	Falko Kuester	Joint Research	3D from video	LDRD	Mark Duchaineau
UC San Diego	Randy Bank	Consultant	Numerical solutions of partial differential equations, multilevel iterative methods, and adaptive grid generations	ASC	Robert Falgout and Panayot Vassilevski
UC San Diego	Charles Elkan	Collaboration, Lawrence Scholar Program	Learning from presence-only data	UCOP	Tina Eliassi-Rad
UC Santa Cruz	Stan Woosley	Joint Research	Computational Astrophysics Consortium	ASCR SciDAC	Louis Howell
UC Santa Cruz	Ethan Miller and Carlos Maltzahn	Collaboration, Lawrence Scholar Program	Semantic file systems	LDRD	Maya Gokhale
Texas A&M University	Bjarne Stroustrup and Lawrence Rauchwerger	Joint Research	Compiler construction and parallel optimizations	ASCR	Dan Quinlan
Texas A&M University	Raytcho Lazarov	Joint Research	Algebraic multigrid for problems for H(curl)	ASCR Base	Panayot Vassilevski
Texas A&M University	Marv Adams	Joint Research	Numerical methods for radiation transport	ASC	Peter Brown
Texas A&M University	Nancy Amato	Collaboration, Lawrence Scholar Program	Novel mechanisms to understand and improve load balance in MPI applications	UCOP	Bronis de Supinski

Academic Outreach (continued)

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
Texas A&M University	Nancy Amato	Collaboration, Lawrence Scholar Program	Parallel graph algorithms	UCOP	Maya Gokhale
Atomic Weapons Establishment	Matt Wheeler and Paul Selby	Joint Research	Development and use of VisIt on Atomic Weapons Establishment applications	ASC	Mark Duchaineau
Ball State University	Irene Livshits	Subcontract	Algebraic multigrid algorithms for finding many eigenpairs of partial differential operators	ASCR	Robert Falgout
Brigham Young University	Bryan Morse	Subcontract	Mosaics and super-resolution of unmanned aerial vehicle-acquired video using locally adaptive warping	LDRD	Mark Duchaineau and Jon Cohen
California Institute of Technology	Michael Ortiz	ASC Predictive Science Academic Alliance Program Center	Center for the Predictive Modeling and Simulation of High-Energy Density Dynamic Response of Materials	ASC	Dick Watson
Cambridge University	Nikos Nikiforkis	Joint Research	Simulation and modeling using Overture	ASCR Base	Bill Henshaw
Carnegie Mellon University	Christos Faloutsos	Subcontract	Mining large dynamic weighted graphs	LDRD	Tina Eliassi-Rad
Carnegie Mellon University	Hanghang Tong	Collaboration	Mining large dynamic weighted graphs	LDRD	Tina Eliassi-Rad
Colorado State University	Donald Estep	Subcontract	A priori and a posteriori analysis of smoothed particle hydrodynamics-finite element methods for pore-scale simulations	LDRD	Carol Woodward
Colorado State University	Michelle Strout and Sanjay Rajopadhye	Collaboration	Program analysis	ASCR	Dan Quinlan
Columbia University	David Keyes	Consultant	Algorithms for the solution of partial differential equations on massively parallel computers; general high-performance computing	LLNL Overhead	James McGraw
Columbia University	David Keyes	Joint Research	Towards optimal petascale simulations	ASCR SciDAC	Robert Falgout
Columbia University	Ian Liplin	Joint Research	Viral discovery and microarrays	DTRA	Tom Slezak
Commissariat à l'Énergie Atomique	Thierry Carrard	Joint Research	Multiresolution visualization techniques	ASC	Mark Duchaineau

Academic Outreach (continued)

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
Chalmers University, Sweden	Sally McKee	Collaboration	Leveraging OpenAnalysis for alias analysis within ROSE	ASC	Dan Quinlan
Cornell University	Claire Cardie	Joint Research	Coreference resolution	DHS	David Buttler
Duke University	Herbert Edelsbrunner	Joint Research	Discrete methods for computing continuous functions	LDRD	Peer-Timo Bremer
Électricité de France	Phillippe Lafon	Collaboration	Aeroacoustics		Bill Henshaw
Imperial College	Paul Kelly and Jose Gabriel de Figueiredo Coutinho	Collaboration	Field-programmable gate arrays research	ASCR	Dan Quinlan
Indiana University	Jeremiah Wilcock	Joint Research	Binary analysis	ASCR	Dan Quinlan
Johns Hopkins University	Allan Boyles	Collaboration	Seismoacoustic modeling for defense-related efforts		Shawn Larsen
Krell Institute	Lucille Kilmer	Subcontract	Department of Energy High-Performance Computer Science Fellowship Program: LLNL portion	ASC	John May
North Carolina State University	Frank Mueller	Joint Research	Compressing message-passing interface traces	ASCR, ASC	Bronis de Supinski
North Carolina State University	Vincent Freeh	Joint Research	Power-aware computing for message-passing interface programs	ASCR, ASC	Bronis de Supinski
Northern Arizona University	Paul Kleim	Joint Research	Microbial forensics	DHS	Tom Slezak
Ohio State University	P. Sadayappan and Christophe Alias	Collaboration	Optimizing compiler program analysis	ASCR	Dan Quinlan
Ohio State University	Joel Saltz and Tahsin Kurc	Collaboration	Data science algorithms		Ghaleb Abdulla
Penn State University	Jinchao Xu, James Brannick, and Ludmil Zikatanov	Subcontract	Multigrid methods for numerical models arising in plasma simulations and quantum field theories	ASCR	Robert Falgout

Academic Outreach (continued)

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
Pennsylvania State University	Ludmil Zikatonov and Robert Schechl	Subcontract	Construction of coarse scale numerical models based on energy minimization and applications to upscaling techniques	ASCR Base	Panayot Vassilevski
Portland State University	Karen Karavanic	Joint Research	Performance analysis infrastructure for petascale applications	ASC CSSE	John May
Princeton Plasma Physics Laboratory	Ravi Samtaney	Collaboration	Implicit methods for magnetohydrodynamics	ASCR SciDAC	Lori Diachin
Princeton University	Adam Burrows and Jason Nordhaus	Joint Research	Computational Astrophysics Consortium	ASCR SciDAC	Louis Howell
Purdue University	Saurabh Bagchi	Joint Research	Distributed computing systems and cyber security	LDRD	Celeste Matarazzo
Purdue University	Ziqiang Cai	Summer Faculty	A posteriori error estimates for partial differential equations	ASC	Robert Falgout
Purdue University	Jayathi Murthy	ASC Predictive Science Academic Alliance Program Center	Center for Prediction of Reliability, Integrity, and Survivability of Microsystems (PRISM)	ASC	Dick Watson
Rensselaer Polytechnic Institute	Don Schwendeman	Subcontract	Development of numerical methods for mathematical models of high-speed reactive and nonreactive flow	ASCR Base	Bill Henshaw
Rensselaer Polytechnic Institute	Ken Jansen	Joint Research	Interoperable technologies for advanced petascale simulation	ASCR SciDAC	Lori Diachin
Rensselaer Polytechnic Institute	Mark Shephard	Joint Research	Terascale simulation tools and technologies	ASCR SciDAC	Lori Diachin
Rice University	John Mellor-Crummey	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
Rice University	John Mellor-Crummey, Keith Cooper, and Vivek Sarkar	Collaboration	Use of ROSE for compiler optimizations	ASCR	Dan Quinlan
Royal Institute of Technology, Sweden	Heinz-Otto Kreiss	Consultant	Adaptive methods for partial differential equations	ASCR Base	Lori Diachin and Anders Petersson

Academic Outreach (continued)

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
Southern Methodist University	Thomas Hagstrom	Joint Research	High-order structure grid methods for wave propagation on complex unbounded domains	ASCR Base	Bill Henshaw
Southern Methodist University	Dan Reynolds	Joint Research	Implicit solvers and preconditioning techniques for simulations of magnetohydrodynamics and core-collapse supernovae	ASCR Base	Carol Woodward
Stanford University	Parvis Moin	ASC Predictive Science Academic Alliance Program Center	Center for Predictive Simulations of Multiphysics Flow Phenomena with Application to Integrated Hypersonic Systems	ASC	Dick Watson
Stanford University	Olav Lindtjorn	Collaboration	Simulation of historical and scenario earthquakes in California		Shawn Larsen
Stanford University/ SLAC National Accelerator Laboratory	Jack Becla and Daniel Wang	Joint Research	Large-scale data management and Large Synoptic Survey Telescope	LDRD	Celeste Matarazzo
State University of New York, Stony Brook	Doug Swesty	Joint Research	Solvers for supernova simulation	ASCR SciDAC	Carol Woodward
State University of New York, Stony Brook	Jim Glimm	Joint Research	Terascale simulation tools and technologies	ASCR SciDAC	Lori Diachin
State University of New York, Stony Brook	Xiao-Lin Li	Joint Research	Interoperable technologies for advanced petascale simulation	ASCR SciDAC	Lori Diachin
State University of New York, Stony Brook	Michael Zingale and Douglas Swesty	Joint Research	Verification and validation of radiation hydrodynamics for astrophysical applications	LDRD	Louis Howell
Swiss Federal Institute of Technology	Stephan Brunner	Subcontract	Efficient numerical algorithms for Vlasov simulation of laser–plasma interactions	LDRD	Jeff Hittinger
Technical University of Vienna	Markus Schordan	Collaboration	Compiler construction	ASCR	Dan Quinlan

Academic Outreach (continued)

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
Tufts University	Scott MacLachlan	Joint Research	Algebraic multigrid algorithms	ASC	Robert Falgout
University of British Columbia	Carl Olivier-Gooch	Subcontract	Interoperable technologies for advanced petascale simulation	ASCR SciDAC	Lori Diachin
University of Chicago	Don Lamb	ASC Academic Strategic Alliance Program Center	Center for Astrophysical Thermonuclear Flashes	ASC	Dick Watson
University of Colorado	Tom Manteuffel	Joint Research	Solution methods for transport problems	ASC	Peter Brown
University of Colorado	Steve McCormick	Joint Research	Towards optimal petascale simulations	ASCR SciDAC	Robert Falgout
University of Colorado	Steve McCormick, Tom Manteuffel, John Ruge, Marian Brezina, Minho Park, Geoff Sanders, and Christian Ketelsen	Subcontract	Geometric and algebraic multigrid methods for quantum chromodynamics, magnetohydrodynamics, elasticity, transport, and other DOE applications	ASCR, ASC	Robert Falgout
University of Colorado	Tom Manteuffel	Joint Research	Towards optimal petascale simulations	ASCR SciDAC	Robert Falgout
University of Colorado	John Ruge	Joint Research	Multipass interpolation for algebraic multigrid	ASC	Ulrike Yang
University of Delaware	Richard Braun	Collaboration	Models of the eye	ASCR Base	Bill Henshaw
University of Delaware	Calvin Keeler	Subcontract	Avian influenza resequencing microarray		Tom Slezak
University of Georgia	David Lowenthal	Joint Research	Power-aware computing for message-passing interface programs; scalable performance models	ASCR, ASC	Bronis de Supinski
University of Houston	Yuriy Fofanov	Joint Research	Genomic algorithms	DHS	Tom Slezak
University of Illinois, Urbana-Champaign	Michael Heath	ASC Academic Strategic Alliance Program Center	Center for Simulation of Advanced Rockets	ASC	Dick Watson

Academic Outreach (continued)

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
University of Illinois/National Center for Supercomputing Applications	Dave Semeraro	Joint Research	VisIt technology applications on National Center of Supercomputing Applications Blue Water system	ASC	Mark Duchaineau
University of Louisville	Yongsheng Liam	Collaboration	Micro-air vehicles		Bill Henshaw
University of Maryland	Jeff Hollingsworth	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
University of Massachusetts Amherst	Andrew McCallum	Joint Research	Cross-language topic models	LDRD	David Buttler
University of Michigan	R. Paul Drake	ASC Predictive Science Academic Alliance Program Center	Center for Radiative Shock Hydrodynamics (CRASH)	ASC	Dick Watson
University of Munich	Dieter August Kranzlmueller	Joint Research	Detecting communication patterns to optimize applications	ASCR, ASC	Bronis de Supinski
University of Nevada, Reno	John Louie	Collaboration	Seismic modeling in the basin and range region		Shawn Larsen
University of North Carolina	Robert Fowler	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
University of Pittsburgh	Janyce Wiebe	Joint Research	DHS: University Affiliates Centers	DHS	David Buttler
University of Southern California	Robert Lucas, Mary Hall, and Jacqueline Chame	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
University of Southern California	Gerard Medioni	Subcontract	Activity analysis in wide-area overhead video	DOE Nonproliferation	Mark Duchaineau
University of Tennessee	Jack Dongarra and Shirley Moore	Joint Research	Performance Engineering Research Institute	ASCR SciDAC	Bronis de Supinski
University of Tennessee	Jack Dongarra	Joint Research	Empirical tuning	ASCR	Dan Quinlan

Academic Outreach (continued)

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
University of Texas, Austin	Robert Moser	ASC Predictive Science Academic Alliance Program Center	Center for Predictive Engineering and Computational Sciences (PECOS)	ASC	Dick Watson
University of Texas, Austin	Omar Ghattas	Joint Research	Towards optimal petascale simulations	ASCR SciDAC	Robert Falgout
University of Texas, San Antonio	Qing Yi	Subcontract	Program analysis and optimization for the empirical tuning of scientific applications	ASCR	Dan Quinlan
University of Utah	Gannesh Gopalakrishnan	Collaboration	Message-passing interface optimizations	ASCR	Dan Quinlan
University of Utah	Ellen Riloff	Subcontract	Coreference resolution research	LDRD	David Buttler
University of Utah	David Pershing	ASC Academic Strategic Alliance Program Center	Center for the Simulation of Accidental Fires and Explosions	ASC	Dick Watson
University of Utah	Chris Johnson, Valerio Pascucci, Chuck Hansen, Claudio Silva, Lee Myers, Allen Sanderson, and Steve Parker	Joint Research	Visualization and Analytics Center for Enabling Technologies	ASCR SciDAC	Mark Duchaineau
University of Utah	Charles Hansen	Consultant	Data exploration, multiresolution scientific data visualization, and algorithm design	ASC CSSE	Mark Duchaineau
University of Utah	Claudio Silva	Joint Research	Studying the topology of point-set surfaces	LDRD	Peer-Timo Bremer
University of Utah	Valerio Pascucci	Subcontract	Algorithms for geometric processing, image segmentation, and data streaming	LDRD	Peer-Timo Bremer
University of Washington	Carl Ebeling	Subcontract	Coarse-grain processor architectures		Maya Gokhale
University of Waterloo	Hans de Sterck	Joint Research	Long-distance interpolation for algebraic multigrid	ASCR SciDAC	Ulrike Yang
University of Wisconsin	Ben Liblit	Joint Research	Scalable debugging	ASCR, ASC	Bronis de Supinski
University of Wisconsin	Jason Kraftcheck	Joint Research	Mesquite software development	ASCR Base	Lori Diachin

Academic Outreach (continued)

University	Faculty	Activity Type	Topic	Funding Source	LLNL Contact
University of Wisconsin	Bart Miller	Joint Research	Support for enhanced Dyninst testing and initial steps toward open binary editing environment	ASC CSSE	Martin Schulz
University of Wuppertal	Karsten Kahl	Visiting Researcher	Algebraic multigrid algorithms	ASC	Robert Falgout
University of Wyoming	Victor Ginting	Subcontract	A priori and a posteriori analysis of smoothed particle hydrodynamics-finite element methods for pore-scale simulations	LDRD	Carol Woodward
Virginia Institute of Technology	Demitrios Nikolopoulos	Joint Research	Power-aware computing for OpenMP programs	ASCR, ASC	Bronis de Supinski
Virginia Institute of Technology	Kirk Cameron	Joint Research	Power-aware computing for hybrid systems	ASCR, ASC	Bronis de Supinski

Publications

Books and Book Chapters

Gregorski, B., J. Senecal, M. Duchaineau, and K. I. Joy, "Compression and Occlusion Culling for Fast Isosurface Extraction from Massive Datasets," in *Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration*, T. Möller, B. Hamann, and R. D. Russell (Springer, Berlin).

Kamath, C., *Scientific Data Mining: A Practical Perspective* (SIAM), 286 pgs.

Labov, S. and T. R. Slezak, "The Indispensable Technology: Detectors for Nuclear, Biological, and Chemical WMD," in *WMD Terrorism: Science and Policy Choices*, S. M. Maurer (MIT Press), pp. 139–168.

Journal Papers

Allen, J. E., S. N. Gardner, E. A. Vitalis, and T. R. Slezak, "Conserved Amino Acid Markers from Past Influenza Pandemic Strains," *BMC Microbiol.* **9**, (LLNL-JRNL-409735).

Amendt, P. A., et al. "Electric Field and Ionization-Gradient Effects on Inertial-Confinement-Fusion Implosions," *Plasma Phys. Contr. F.* **51** (12).

Anderson, J. C., L. J. Gosink, M. A. Duchaineau, and K. I. Joy, "Interactive Visualization of Function Fields by Range–Space Segmentation," *Comput. Graph. Forum* **28** (3), 727–734.

Anenberg, S. C., et al. "Intercontinental Impacts of Ozone Pollution on Human Mortality," *Environ. Sci. Technol.* **43** (17), 6482–6487.

Anisimov, A. P., et al. "Amino Acid and Structural Variability of *Yersinia pestis* LcrV Protein," *Infection, Genetics, and Evolution* **10** (1), 137–145 (LLNL-JRNL-420969).

Appelo, D. and N. A. Petersson, "A Stable Finite Difference Method for the Elastic Wave Equation on Complex Geometries with Free Surfaces," *Commun. Comput. Phys.* **5** (1), 84–107 (UCRL-JRNL-237329).

Baker, A. H., E. R. Jessup, and T. R. Kolev, "A Simple Strategy for Varying the Restart Parameter in GMRES(m)," *J. Comput. Appl. Math.* **230** (2), 751–761 (UCRL-JRNL-235266).

Baker, A. H., T. R. Kolev, and U. M. Yang, "Improving Algebraic Multigrid Interpolation for Linear Elasticity Problems," *Numer. Linear Algebr.* (LLNL-JRNL-412928).

Banks, J. W. and J. N. Shadid, "An Euler System Source Term that Develops Prototype Z-Pinch Implosions Intended for the Evaluation of Shock-Hydro Methods," *Int. J. Numer. Meth. Fl.* **61** (7), 725–751.

Banks, J. W., W. D. Henshaw, and J. N. Shadid, "An Evaluation of the FCT Method for High-Speed Flows on Structured Overlapping Grids," *J. Comput. Phys.* **228** (15), 5349–5369.

Bihari, B. L. and P. N. Brown, "A Linear Algebraic Analysis of Diffusion Synthetic Acceleration for the Boltzmann Transport Equation II: The Simple Corner Balance Method," *SIAM J. Numer. Anal.* **47** (3), 1782–1826.

Bronevetsky, G., J. Gyllenhaal, and B. R. de Supinski, "CLOMP: Accurately Characterizing OpenMP Application Overheads," *Int. J. Parallel Prog.* **37** (3), 250–265 (LLNL-JRNL-408738).

Chakicherla, A., et al. "A SpaK/SpaR Two-Component System Characterized by a Structure-Driven Domain-Fusion Method and In Vitro Phosphorylation Studies," *PLoS Comput. Biol.* **5** (6), (LLNL-JRNL-411091).

Chen, J. H., et al. "Terascale Direct Numerical Simulations of Turbulent Combustion Using S3D," *Comput. Sci. Eng.* **2**, (LLNL-JRNL-409355).

Das, T., et al. "Structure and Detectability of Trends in Hydrological Measures Over the Western United States," *J. Hydrometeorol.* **10** (4), 871–892 (LLNL-JRNL-406809).

Day, M., et al. "Turbulence Effects on Cellular Burning Structures in Lean Premixed Hydrogen Flames," *Combust. Flame* **156** (5), 1035–1045.

Fattebert, J.-L., et al. "Quantitative Assessment of Electrostatic Embedding in Density Functional Theory Calculations of Biomolecular Systems," *J. Chem. Theory Comput.* **5** (9), 2257–2264.

Fiore, A. M., et al. "Multimodel Estimates of Intercontinental Source-Receptor Relationships for Ozone Pollution," *J. Geophys. Res.-Atmos.* **114** (D04301), (LLNL-JRNL-407973).

Gardner, S. N., et al. "Multiplex Primer Prediction Software for Divergent Targets," *Nucleic Acids Res.* **37** (19), 6291–6304.

Gates, D. A., et al. "Overview of Results from the National Spherical Torus Experiment (NSTX)," *Nucl. Fusion* **49** (10), (LLNL-JRNL-409680).

Gibson, C. R., et al. "Design of the NIF Cryogenic Target System," *Fusion Sci. Technol.* **55** (3), 233–236 (LLNL-JRNL-404974).

Goltsman, D. S. A., et al. "Community Genomic and Proteomic Analyses of Chemoautotrophic Iron-Oxidizing *Leptospirillum rubarum* (Group II) and *Leptospirillum ferrodiazotrophum* (Group III) Bacteria in Acid Mine Drainage Biofilms," *Appl. Environ. Microb.* **75** (13), 4599–4615 (LLNL-JRNL-421044).

Henshaw, W. D. and K. K. Chand, "A Composite Grid Solver for Conjugate Heat Transfer in Fluid-Structure Systems," *J. Comput. Phys.* **228** (10), 3708–3741 (LLNL-JRNL-407079).

Hidalgo, H. G., et al. "Detection and Attribution of Streamflow Timing Changes to Climate Change in the Western United States," *J. Climate* **22** (13), 3838–3855 (UCRL-JRNL-236583).

Keedy, D. A., et al. "The Other 90% of the Protein: Assessment Beyond the Calphas for CASP8 Template-Based and High-Accuracy Models," *Proteins: Structure, Function, and Bioinformatics* **77** (S9), 29–49 (LLNL-JRNL-420970).

Kirkwood, R. K., et al. "Sensitivity of Ignition-Scale Backlit Thin-Shell Implosions to Hohlraum Symmetry in the Foot of the Drive Pulse," *Phys. Plasmas* **16** (1), (LLNL-JRNL-406419).

Klein, S. A., et al. "Intercomparison of Model Simulations of Mixed-Phase Clouds Observed During the ARM Mixed-Phase Arctic Cloud Experiment. I: Single-Layer Cloud," *Q. J. Roy. Meteor. Soc.* **135** (641), 979–1002 (LLNL-JRNL-401869).

Kolev, T. V. and P. S. Vassilevski, "Parallel Auxiliary Space AMG for H(curl) Problems," *J. Comput. Math.* **27** (5), 604–623 (UCRL-JRNL-237306).

Kolev, T. V. and R. N. Rieben, "A Tensor Artificial Viscosity Using a Finite Element Approach," *J. Comput. Phys.* **228** (22), 8336–8366.

Kupiainen, M. and B. Sjögreen, "A Cartesian Embedded Boundary Method for the Compressible Navier–Stokes Equations," *J. Sci. Comput.* **41** (1), 94–117 (LLNL-JRNL-402504).

Lobell, D., et al. "Regional Differences in the Influence of Irrigation on Climate," *J. Climate* **22** (8), 2248–2255 (LLNL-JRNL-401185).

Modjaz, M., et al. "From Shock Breakout to Peak and Beyond: Extensive Panchromatic Observations of the Type Ib Supernova 2008D Associated with Swift X-Ray Transient 080109," *Astrophys. J.* **702** (1), 226–248.

Morrison, H., et al. "Intercomparison of Model Simulations of Mixed-Phase Clouds Observed During the ARM Mixed-Phase Arctic Cloud Experiment. II: Multilayer Cloud," *Q. J. Roy. Meteor. Soc.* **135** (641), 1003–1019 (LLNL-JRNL-401832).

Murray, C., "Accelerating into the Future," *Phys. World* **22** (5), 16–17.

Noeth, M., et al. "ScalaTrace: Scalable Compression and Replay of Communication Traces for High-Performance Computing," *J. Parall. Distr. Com.* **69** (8), 696–710 (LLNL-JRNL-403992).

Ogitsu, T., et al. "First-Principles Calculations of Solid and Liquid Aluminum Optical Absorption Spectra Near the Melting Curve: Ambient and High-Pressure Results," *Phys. Rev. B* **80** (21).

Petersson, N. A. and B. Sjögreen, "An Energy Absorbing Far-Field Boundary Condition for the Elastic Wave Equation," *Commun. Comput. Phys.* **6** (3), 483–508 (LLNL-JRNL-405423).

Reidmiller, D. R., et al. "The Influence of Foreign vs. North American Emissions on Surface Ozone in the U.S.," *Atmos. Chem. Phys.* **9** (14), 5027–5042.

Robey, H. F., et al. "Hohlraum-Driven Mid-Z (SiO_2) Double-Shell Implosions on the Omega Laser Facility and their Scaling to NIF," *Phys. Rev. Lett.* **103** (14).

Seager, M. and B. Gorda, "A Collaboration and Commercialization Model for Exascale Software Research," *Int. J. High Perform. C.* **23** (4), 395–397 (LLNL-JRNL-414976).

Seager, M. and B. Gorda, "The Case for a Hierarchical System Model for Linux Clusters," *Int. J. High Perform. C.* **23** (4), 350–354 (LLNL-JRNL-413749).

Singer, M. A. and W. H. Green, "Using Adaptive Proper Orthogonal Decomposition to Solve the Reaction-Diffusion Equation," *Appl. Numer. Math.* **59** (2), 272–279 (UCRL-JRNL-236975).

Singer, M. A., W. D. Henshaw, and S. L. Wang, "Computational Modeling of Blood Flow in the TrapEase Inferior Vena Cava Filter," *J. Vasc. Interven. Radiol.* **20** (6), 799–805 (LLNL-JRNL-401016).

Sjögreen, B. and H. C. Yee, "Variable High-Order Multiblock Overlapping Grid Methods for Mixed Steady and Unsteady Multiscale Viscous Flows," *Commun. Comput. Phys.* **5** (2–4), 730–744.

Temple, G., et al. "The Completion of the Mammalian Gene Collection (MGC)," *Genome Res.* **19** (12), 2324–2333.

Vishnu, A., et al. "Topology Agnostic Hot-Spot Avoidance with InfiniBand," *Concurr. Comp.-Pract. E.* **21** (3), 301–319.

Wang, W., C. W. Shu, H. C. Yee, and B. Sjögreen, "High-Order Well-Balanced Schemes and Applications to Non-Equilibrium Flow with Stiff Source Terms," *J. Comput. Phys.* **228** (18), 6682–6702 (LLNL-JRNL-409903).

Williams, D. N., et al. "The Earth System Grid: Enabling Access to Multimodel Climate Simulation Data," *B. Am. Meteorol. Soc.* **90** (2).

Wilson, W. C., et al. "A Multiplex Real-Time Reverse Transcription Polymerase Chain Reaction Assay for Detection and Differentiation of Bluetongue Virus and Epizootic Hemorrhagic Disease Virus Serogroups," *J. Vet. Diagn. Invest.* **21** (6), 760–770.

Wu, D., et al. "A Phylogeny-Driven Genomic Encyclopedia of Bacteria and Archaea," *Nature* **462**, 1056–1060.

Xu, X. Q., et al. "Dynamics of Kinetic Geodesic-Acoustic Modes and the Radial Electric Field in Tokamak Neoclassical Plasmas," *Nucl. Fusion* **49** (6), (LLNL-JRNL-409472).

Yang, U. M., "On Long-Range Interpolation Operators for Aggressive Coarsening," *Numer. Linear Algebr.* (LLNL-JRNL-417371).

Yee, H. C. and B. Sjögreen, "Performance of High-Order Filter Methods for a Richtmyer–Meshkov Instability," *Shock Waves*, (Springer), **17**, 185–193.

Conference Proceedings

Amendt, P., C. Cerjan, J. Milovich, and H. Robey, "Prospects for Demonstrating Double-Shell Ignition on the National Ignition Facility," *Current Trends in International Fusion Research: Proceedings of the 7th Symposium—Selected Presentations*, (Amer. Inst. Phys.), **1154**, 234–234.

Bronevetsky, G., J. Gyllenhaal, and B. R. de Supinski, "CLOMP: Accurately Characterizing OpenMP Application Overheads," *International Workshop on OpenMP*, (Springer/Plenum), **37**, 250–265 (LLNL-CONF-401267).

Chuang, C., D. Bergmann, and P. Cameron-Smith, "Global Modeling of Atmospheric Aerosols Including SOA," *DOE Atmospheric Science Program Science Team Meeting*

Dahlgren, T. L. and I. Abramova, "Toward Multilanguage, Multicomponent Interface Contract Enforcement," *Proceedings of the 2009 Workshop on Component-Based High Performance Computing*, (ACM).

Damevski, K., H. Chen, and T. L. Dahlgren, "Reducing Component Contract Overhead by Offloading Enforcement," *Proceedings of the 2009 Workshop on Component-Based High Performance Computing*, (ACM).

de Supinski, B. R., et al. "Modeling the Office of Science 10-Year Facilities Plan: The PERI Architecture Tiger Team," *SciDAC 2009* (J. Phys.: Conf. Ser., IOP Publishing), **180**.

Fowler, R., et al. "Frontiers of Performance Analysis on Leadership Class Systems," *SciDAC 2009* (J. Phys.: Conf. Ser., IOP Publishing), **180**.

Henshaw, W. D., "Solving Partial Differential Equations on Overlapping Grids," *ASTRONUM 2008*, (Astronomical Society Pacific), **406**, 231–236 (LLNL-PROC-407183).

Hilbrich, T., B. R. de Supinski, M. Schulz, and M. Mueller, "A Graph-Based Approach for MPI Deadlock Detection," *Proceedings of the 23rd International Conference on Supercomputing*, (ACM), 293–305.

Kumar, V., et al. "Architectural Implications for Spatial Object Association Algorithms," *IEEE International Symposium on Parallel and Distributed Processing*, (IEEE), (LLNL-CONF-410283).

Li, J., et al. "Machine-Learning-Based Online Performance Prediction for Runtime Parallelization and Task Scheduling," *IEEE International Symposium on Performance Analysis of Systems and Software*, (IEEE), (LLNL-CONF-407723).

Li, L., T. L. Dahlgren, L. C. McInnes, and B. I. Norris, "Interface Contract Enforcement for Improvement of Computational Quality of Service (CQoS) for Scientific Components," *Workshop on Component-Based High Performance Computing*, (ACM).

Liao, C. H., D. J. Quinlan, J. J. Willcock, and T. Panas, "Extending Automatic Parallelization to Optimize High-Level Abstractions for Multicore," *International Workshop on OpenMP*, (Lect. Notes Comput. Sc., Springer–Verlag Berlin), **5568**, 28–41 (LLNL-CONF-409406).

Rountree, B., et al. "Adagio: Making DVS Practical for Complex HPC Applications," *23rd International Conference on Supercomputing*, (ACM), 460–469.

Schulz, M., et al. "On the Performance of the Miranda CFD Code on Multicore Architectures," *21st International Conference on Parallel Computational Fluid Dynamics (ParallelCFD 2009)*.

RETURN TO TABLE OF CONTENTS

Acronyms

ALE ....Arbitrary Lagrangian–Eularian
AMR.................adaptive mesh refinement
ASC ....Advanced Simulation and Computing
BCB ....BlueCoat Blocking
CASC ....Center for Applied Scientific Computing
CASTLE ....Collaborative Authorization for the Safety-Basis Total Lifecycle Environment
CIO ....Chief Information Office
CMDB ....Configuration Management Database
CSSM ....Cyber Security Site Manager
DoD ....Department of Defense
DOE ....Department of Energy
ExaCT ....Exascale Computing Technologies
FN ....foreign national
GB....................gigabyte
HPC ....high-performance computing
I-BC ....Instrument-Based Controller
ICCS ....Integrated Computer Control System
IM ....instant messaging
I/O...................input and output
IP ....Internet Protocol
ISCR ....Institute for Scientific Computing Research
IT ....information technology
LDRD ....Laboratory Directed Research and Development
LLNL................Lawrence Livermore National Laboratory
LOFT ....Laptops on Foreign Travel
LPI...................laser–plasma interaction
LSO ....Livermore Site Office
MPI..................message passing interface
NAP..................National Nuclear Security Administration Policy Letters
NEM.................Nuclear Enterprise Modeling
NIF ....National Ignition Facility
NNSA ....National Nuclear Security Administration
petaFLOPS ......quadrillion floating-point operations per second
PHG ....Parallel Hypergraph
PMESH.............Parallel Mesh Generation
RCB ....recursive coordinate bisection
S&T ....Science and Technology
SAN ....storage area network
SSCL ....Site Security Component Library
SCR ....Scalable Checkpoint/Restart
SETAC.............Supercomputing Enabled Transformational Analytics Capability
SISC ....Storage Intensive Supercomputing
SLURM.............Simple Linux Utility for Resource Management
SOA ....Service Oriented Architecture
SPLAT ....Shot Planning and Analysis Tool
SSA...................Space Situational Awareness
SSCM ....Site Security Component Library Manager
SWL ....Synthetic WorkLoad
teraFLOPS ....trillion floating-point operations per second
TESSA ....Testbed Environment for Space Situational Awareness
TLCC ....Tri-Laboratory Capacity Cluster
TSF ....Terascale Simulation Facility
UQ....................uncertainty quantification
VLAN ....Virtual Local Area Network